

RANDOM LATTICE TRIANGULATIONS: STRUCTURE AND ALGORITHMS

PIETRO CAPUTO, FABIO MARTINELLI, ALISTAIR SINCLAIR,
AND ALEXANDRE STAUFFER

ABSTRACT. The paper concerns *lattice triangulations*, i.e., triangulations of the integer points in a polygon in \mathbb{R}^2 whose vertices are also integer points. Lattice triangulations have been studied extensively both as geometric objects in their own right and by virtue of applications in algebraic geometry. Our focus is on random triangulations in which a triangulation σ has weight $\lambda^{|\sigma|}$, where λ is a positive real parameter and $|\sigma|$ is the total length of the edges in σ . Empirically, this model exhibits a “phase transition” at $\lambda = 1$ (corresponding to the uniform distribution): for $\lambda < 1$ distant edges behave essentially independently, while for $\lambda > 1$ very large regions of aligned edges appear. We substantiate this picture as follows. For $\lambda < 1$ sufficiently small, we show that correlations between edges decay exponentially with distance (suitably defined), and also that the *Glauber dynamics* (a local Markov chain based on flipping edges) is rapidly mixing (in time polynomial in the number of edges in the triangulation). This dynamics has been proposed by several authors as an algorithm for generating random triangulations. By contrast, for $\lambda > 1$ we show that the mixing time is exponential. These are apparently the first rigorous quantitative results on the structure and dynamics of random lattice triangulations.

1. INTRODUCTION

1.1. Background. Let $\Lambda_{m,n}^0 = \{0, 1, \dots, m\} \times \{0, 1, \dots, n\}$ denote the set of integer points in an $m \times n$ rectangle in \mathbb{R}^2 . This paper is concerned with *(full) triangulations* of $\Lambda_{m,n}^0$, i.e., triangulations that use all the points. Any such triangulation partitions the rectangle into $2mn$ triangles each of which is *unimodular* (has area $1/2$). See Figure 1 for an example of a triangulation with $m = 5, n = 7$. Most of our results extend to the case of triangulations of the integer points in an arbitrary (not necessarily convex) lattice polygon.

Lattice triangulations are fundamental discrete geometric objects that have a rich and beautiful structure (see, e.g., [3] for background). For example, a number of recent papers have developed ingenious combinatorial arguments to estimate the asymptotic number of triangulations [1, 8, 13, 16], as well as to explore their connectivity properties under natural local moves [5, 8].

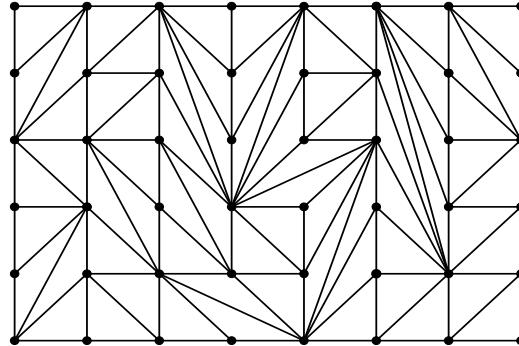
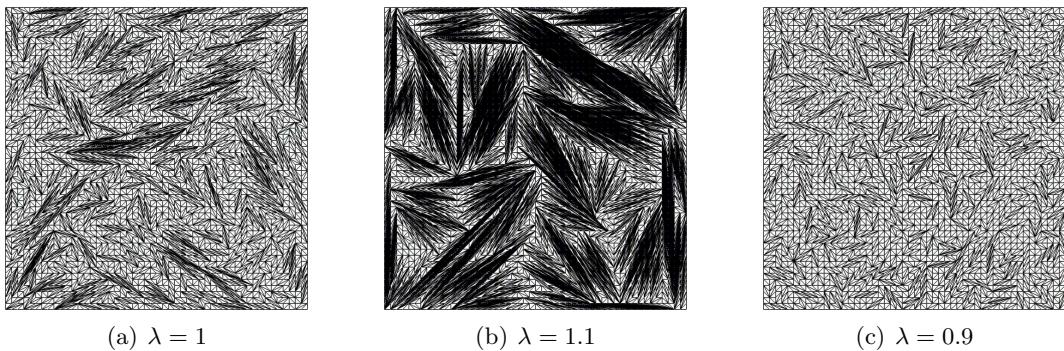
Lattice triangulations have also received much attention in algebraic geometry, through connections with plane algebraic curves and Hilbert’s 16th problem [21], the theory of discriminants [6], and toric varieties [2]. In several of these contexts one is chiefly interested in properties of “typical” triangulations, such as whether they are *regular* (in the sense

Date: November 9, 2012.

2010 *Mathematics Subject Classification.* 60K35, 68W20, 05C81.

Key words and phrases. Triangulations, Spatial mixing, Glauber dynamics, Mixing times, Rapid mixing.

The first, second and fourth authors are supported in part by the European Research Council through AdG “PTRELSS” 228032. Part of this work was done while the third author was visiting University of Roma Tre supported in part by the same grant. The third author was also supported by NSF grant CCF-1016896.

FIGURE 1. *A triangulation of a 5×7 region*FIGURE 2. *Random triangulations of a 50×50 lattice*

of being representable by a “nice” lifting function [20, 24]), whether they contain “long, thin” triangles, etc. This leads one to investigate random triangulations drawn from a uniform distribution [8, 22].

A natural generalization that can yield useful insights here is to introduce weights: specifically, we consider the distribution in which a triangulation σ has weight $\lambda^{|\sigma|}$, where $|\sigma|$ is the total ℓ_1 length of the edges in σ and $\lambda > 0$ is a real parameter. The case $\lambda = 1$ is the uniform distribution, while $\lambda < 1$ (respectively, $\lambda > 1$) favors triangulations with shorter (respectively, longer) edges. Figure 2 shows typical triangulations of a 50×50 region for three values of λ . This weighted version corresponds rather closely to a model that has recently been studied in statistical physics [17].

One striking feature of the pictures in Figure 2 is the tendency in the case $\lambda > 1$ for edges to line up in macroscopically large regions of similar slope, while for $\lambda < 1$ these regions disappear. The uniform case $\lambda = 1$ appears to represent a “phase transition” between these two regimes, in which fairly large regions form but are unstable. One goal of this paper is to initiate a rigorous, quantitative study of this phenomenon.

How does one generate a random triangulation such as those shown in Figure 2? The only known feasible approach for large values of m, n is to simulate a local Markov chain, or *Glauber dynamics*, which converges to the desired probability distribution over triangulations. Fortunately there is a very natural dynamics here: pick a random edge of the triangulation and, if it is the diagonal of a convex quadrilateral (which is in fact always

a parallelogram), flip it to the opposite diagonal. The induced graph on the set of triangulations, in which two triangulations are adjacent iff they differ by one edge flip, is known as the “flip graph.” The flip graph is well known to be connected [9] (indeed, quite a bit more is known about its geometry [5, 8]), so the Glauber dynamics converges to the uniform distribution over triangulations. Moreover, a standard “heat-bath” modification of the flip probabilities achieves the weighted distribution $\lambda^{|\sigma|}$ for any desired λ .

This leads to a fundamental question of both algorithmic and structural interest, posed by Emo Welzl and others [22]: what is the *mixing time* of the Glauber dynamics, i.e., the number of random flips until the dynamics is close to the stationary distribution? The pictures in Figure 2, together with the actual Glauber dynamics simulations used to generate them, suggest the following conjecture. When $\lambda < 1$ there are no long-range correlations between edges, and we would expect the mixing time to be small; when $\lambda > 1$, the large “frozen” regions of aligned edges take a very long time to break up, and consequently the mixing time should be large. We may summarize this in:

Conjecture 1.1. (a) *For any fixed $\lambda < 1$, the mixing time of the Glauber dynamics is $\text{poly}(m + n)$ (specifically, $O(mn(m + n))$). (b) For any fixed $\lambda > 1$, the mixing time is exponential in $(m + n)$ (specifically, $\exp(\Omega(mn(m + n)))$).*

The “critical” case $\lambda = 1$ appears to be rather delicate: while large regions tend to form, they tend not to be stable over long time scales, and the behavior of the mixing time is less clear.

A second goal of this paper is to make the first rigorous progress in analyzing the mixing time of these dynamics, and in particular towards a proof of Conjecture 1.1.

Before describing our results, we relate random triangulations to classical spin systems in statistical physics and explain why they present a challenge to existing analysis techniques. It is well known that, in any triangulation of $\Lambda_{m,n}^0$, the *midpoints* of all the edges are fixed, and are precisely the points of the half-integer lattice (with the original lattice points removed); see the example in Figure 1. Thus we can view any triangulation as an assignment of “spin values” (edges) to each of these midpoints, subject of course to consistency constraints. One might hope, then, to capitalize on the vast literature on both the structure and dynamics of spin systems. However, there is a crucial difference here. In a classical spin system on a graph, all interactions are local; i.e., the distribution of the value of a given spin is determined by the spins of its neighbors in the graph. For triangulations, on the other hand, each edge belongs to two triangles and thus has four neighboring edges—but the midpoints of these edges depend on the triangulation and may lie very far from the midpoint of the edge itself. In other words, the geometry of a triangulation seems to have little to do with the Cartesian geometry of the lattice. It is this non-locality that makes the application of established techniques very challenging here.

1.2. Contributions. Our first results confirm the empirical observation that, in the regime $\lambda < 1$, correlations between edges decay rapidly with distance—a property often referred to in the spin systems literature as “strong spatial mixing.” Specifically, for any sufficiently small $\lambda < 1$, we show that if the configuration of the edge at midpoint x is fixed to σ_x , then the effect on the distribution of the edges at other midpoints $z \in A$, for an arbitrary subset A , decays exponentially with a natural distance $d(A, \sigma_x)$ between A and σ_x . (See Section 4 for the definition of $d(\cdot, \cdot)$.)

Theorem 1.2. *There exists $\lambda_0 > 0$ such that for all $\lambda < \lambda_0$ the following holds. Let μ denote the λ -weighted distribution on triangulations of $\Lambda_{m,n}^0$, and μ^{σ_x} the distribution conditional on the edge at x having configuration σ_x . Then the variation distance between μ and μ^{σ_x} at any set of midpoints A satisfies $\|\mu - \mu^{\sigma_x}\|_A \leq |A| \exp(-c d(A, \sigma_x))$, where the constant $c > 0$ depends only on λ .*

We remark that this result holds in the presence of arbitrary fixed edges, or *constraints*, and therefore for triangulations of arbitrary lattice polygons. Handling constraints introduces some technical complications into our proofs, and necessitates in particular the understanding of minimum length, or “ground state” triangulations. (In the absence of constraints, ground state triangulations are trivial: all edges are either horizontal, vertical or unit diagonal.) Fortunately, as we show, ground state triangulations with constraints can be constructed greedily, by placing each edge independently in its minimum length configuration consistent with the constraints. We also prove another fundamental property in this regime, namely that the probability of an edge exceeding its ground state length by k decays exponentially with k .

Theorem 1.2 is analogous to spatial mixing results in classical spin systems, with the additional twist of the distance function d whose definition is tailored to the geometry of triangulations. Our argument is reminiscent of classical “Peierls-type” arguments for the Ising model, which explains why we require λ to be sufficiently small. We conjecture that the theorem holds for all $\lambda < 1$, but handling values of λ close to 1 will require more sophisticated methods.

We then turn to the Glauber dynamics, again for the regime $\lambda < 1$. Here we can show the following.

Theorem 1.3. *There exists $\lambda_1 > 0$ such that for all $\lambda < \lambda_1$ the mixing time of the Glauber dynamics on triangulations of $\Lambda_{m,n}^0$ with parameter λ (and in the presence of arbitrary constraints) is $O(mn(m+n))$.*

Theorem 1.3 goes some way towards proving part (a) of Conjecture 1.1, though it does require λ to be sufficiently small (rather than merely less than 1). Our proof here uses a path coupling argument with a suitably chosen exponential metric on triangulations, and exploits properties of the geometry of the flip graph which we also establish. We also prove that our bound on the mixing time is tight. In the special case of triangulations of a 1-dimensional region $\Lambda_{1,n}^0$, we show that the mixing time is $O(n^2)$ for all $\lambda < 1$. In this case triangulations turn out to be isomorphic to “lattice paths” of length $2n$ starting and ending at 0 with ± 1 increments at each step.

Our final main result concerns the Glauber dynamics in the complementary regime $\lambda > 1$. Here we are able to prove a slightly weaker version of part (b) of Conjecture 1.1.

Theorem 1.4. *For any $\lambda > 1$, the mixing time of the Glauber dynamics on triangulations of $\Lambda_{m,n}^0$ with parameter λ is $\exp(\Omega(m+n))$.*

While this result confirms the conjectured exponential slowdown for all $\lambda > 1$, we believe that the mixing time should be exponential in the maximum total edge length $mn(m+n)$ rather than in the maximum length of a single edge $(m+n)$. Theorem 1.4 is proved by exhibiting an explicit bottleneck in the dynamics, specifically an initial configuration from which it takes a very long time to change the orientation of some long edge. We also prove a stronger lower bound $\exp(\Omega(n^2/m))$ which holds whenever $m \ll \sqrt{n}$.

1.3. Related work. Triangulations of general point sets are a large topic with numerous applications in mathematics and computer science, including combinatorics, optimization,

algebraic geometry, computational geometry and scientific computing. For excellent background the reader is referred to the survey by Lee [10] and the recent book of DeLoera, Rambau and Santos [3], which also discusses lattice triangulations in some depth. As mentioned earlier, lattice triangulations have been studied both in their own right as geometric objects, and in several contexts in algebraic geometry [2, 6, 20, 21, 24].

Much of the work on lattice triangulations has focused on counting them. A sequence of beautiful combinatorial arguments [16, 1, 8, 13] have shown that the number of triangulations of $\Lambda_{m,n}^0$ is at most $O(6.86^{mn})$ [13] and at least $\Omega(4.15^{mn})$ [8]. Although it is not the tightest upper bound known, we briefly mention the elegant result of Anclin [1], who shows that if the edges of a triangulation of $\Lambda_{m,n}^0$ are added one by one starting at the top left and proceeding left-to-right and top-to-bottom (by midpoint), then the maximum number of choices for each edge is two. Since there are fewer than $3mn$ interior edges, this immediately yields an upper bound of 8^{mn} on the number of triangulations. (In contrast, the best known upper bound for the number of triangulations of a general set of n points in the plane is 43^n [18].)

The flip graph on triangulations of $\Lambda_{m,n}^0$ has also been studied in some depth, though we should stress that our work is apparently the first to handle constraints (and thus general lattice polygons). For example, the flip graph is known to have diameter $O(mn(m+n))$ and to be isometrically embeddable into the hypercube [5]. Random walk on this graph has been proposed by several authors as an algorithm for generating random triangulations, and has been used heuristically in the formulation of conjectures regarding typical triangulations [8, 22]. However, nothing is known rigorously about its mixing time. (We note in passing that the analogous flip dynamics for triangulations of (the vertex set of) a convex polygon has been extensively analyzed [14, 15]; however, that problem is much easier because the set of triangulations has a Catalan structure.)

Random lattice triangulations with weights have appeared in slightly different form as a model in statistical physics [17]; there, triangulations are weighted according to the sum of squares of their vertex degrees. This is actually very close to our model, as high degrees are associated with long edges. The structural results reported in [17] are based on simulations and are non-rigorous. Rigorous results have been obtained for a loosely related “topological glass” model by Eckmann and Younan [4].

The literature on structural properties and Glauber dynamics of lattice spin systems is too vast to summarize here. We refer the reader to the standard references [19] for structural properties such as spatial mixing, and [12] for mixing times of Glauber dynamics. As explained earlier, while triangulations may be viewed as a spin system, their geometry is very different from that of a traditional spin system on the lattice; our paper can be seen as a first step towards obtaining structural and mixing time results for triangulations analogous to those for classical spin systems.

We mention finally that our mixing time result for the special case of 1-dimensional regions with $\lambda < 1$ is related to work of Greenberg, Pascoe and Randall [7] on lattice paths. Those authors use a similar path coupling argument, but for a different probability distribution on lattice paths: in their model paths are biased according to the area under the path, while in ours the bias depends on the excursions of the path from a fixed line.

2. THE MODEL

2.1. Lattice triangulations. Let $\Lambda_{m,n}^0$ denote the set of points in the $m \times n$ region of the integer lattice \mathbb{Z}^2 , i.e., $\Lambda_{m,n}^0 := \{0, 1, \dots, m\} \times \{0, 1, \dots, n\}$. A *(full) triangulation*

of $\Lambda_{m,n}^0$ is any maximal set of non-crossing edges (straight line segments), each of which connects two points of $\Lambda_{m,n}^0$ and passes through no other point. We write $\Omega_{m,n}$ for the set of triangulations of $\Lambda_{m,n}^0$. (We will drop the subscripts on Λ^0 and Ω when their values are not important. Also, we will occasionally abuse notation by using $\Lambda_{m,n}^0$ to refer to the geometric region of \mathbb{R}^2 that is the convex hull of the integer points.¹) In this section we spell out a few basic properties of $\Omega_{m,n}$ that are either well known (see, e.g., [3]) or easily deduced.

As for any point set in \mathbb{R}^2 , the numbers of triangles and edges in every triangulation of $\Lambda_{m,n}^0$ are determined by the total number of points and the number of points on the boundary of the convex hull of the point set. Thus every triangulation in $\Omega_{m,n}$ has $(m+1)(n+1)$ points, $2mn$ triangles, and $3mn+m+n$ edges, $2(m+n)$ of which are on the boundary and the remainder in the interior. Moreover, since $\Lambda_{m,n}^0$ consists of the integer points inside a lattice polygon, every triangle is *unimodular*, i.e., has area $\frac{1}{2}$.

In any triangulation, the *midpoints* of the edges are precisely the points of the half-integer lattice with the integer points removed, i.e.,

$$\Lambda_{m,n} := \{0, \frac{1}{2}, 1, \frac{3}{2}, \dots, m - \frac{1}{2}, m\} \times \{0, \frac{1}{2}, 1, \frac{3}{2}, \dots, n - \frac{1}{2}, n\} \setminus \Lambda_{m,n}^0.$$

Thus we may think of a triangulation as an assignment $\sigma = \{\sigma_x\}_{x \in \Lambda_{m,n}}$ of an edge to each point of $\Lambda = \Lambda_{m,n}$. We call midpoints x with one integer and one half-integer coordinate “Type 1”, and those with two half-integer coordinates “Type 2.” Figure 3 shows some of the possible configurations for the edge at Type 1 and Type 2 points x .

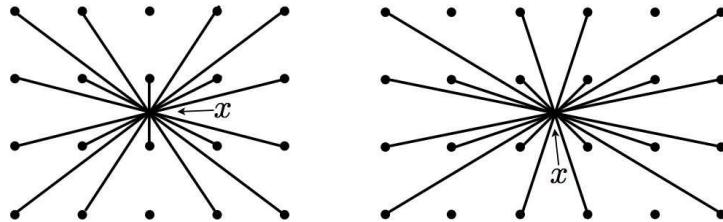


FIGURE 3. Possible configurations of the edge at x . The configurations differ slightly according to whether x is Type 1 (left-hand figure) or Type 2 (right-hand figure). For Type 1 points the minimum length configuration is vertical (or horizontal); for Type 2 it is a unit diagonal.

Note that the minimum length configuration of the edge at a Type 1 point is horizontal or vertical, and at a Type 2 point it is a unit diagonal. The configurations of the boundary edges at points $x = (0, j), (m, j)$ (respectively, $x = (i, 0), (i, n)$) are forced to be horizontal (respectively, vertical) in all triangulations. We call triangulations in which all edges have minimal length *ground state* triangulations. There are exactly 2^{mn} ground state triangulations, which are equivalent up to flipping of unit diagonals.

We will also consider a generalization to triangulations with the configurations of some interior edges fixed. Let Δ be an arbitrary subset of Λ , and $\eta = \{\eta_x\}_{x \in \Delta}$ an arbitrary assignment of edges to the points of Δ that is consistent, i.e., none of the edges cross.

¹For notational convenience we shall adopt the slightly non-standard convention that (i, j) denotes the point in \mathbb{R}^2 with vertical coordinate i and horizontal coordinate j . Thus the region $\Lambda_{m,n}^0$ has vertical and horizontal dimensions m and n respectively.

We denote by $\Omega(\eta, \Delta)$ the set of triangulations $\sigma \in \Omega$ that agree with η on Δ , i.e., all possible completions of the partial triangulation η . Since triangulations are just maximal consistent sets of edges, it is clear that for any η at least one such completion always exists. We refer to the edges in η as *constraints*. See Figure 4.

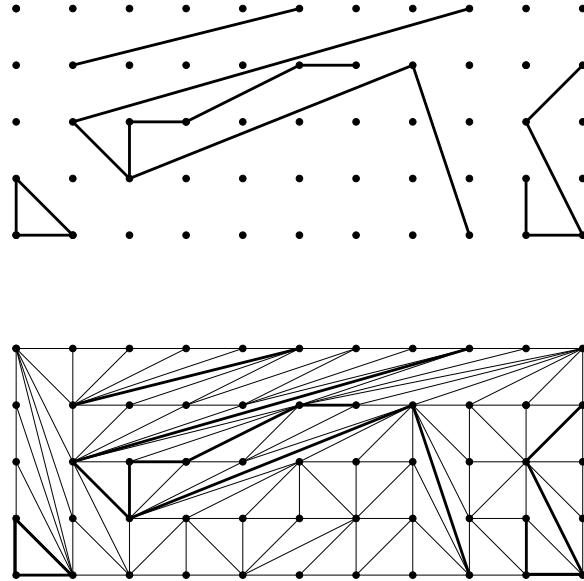


FIGURE 4. *Example of a triangulation of a 4×10 region. Above: constraint edges. Below: a triangulation consistent with the constraints (in bold).*

An important application of constraints is to triangulations of arbitrary lattice polygons (i.e., polygons—not necessarily convex—whose vertices are lattice points in \mathbb{Z}^2). Let P be the set of integer points in a lattice polygon. Then triangulations of the point set P correspond to triangulations of an enclosing rectangle $\Lambda_{m,n}^0$ with the polygon edges as constraints (and an arbitrary fixed triangulation outside the polygon). See Figure 5 for an example.

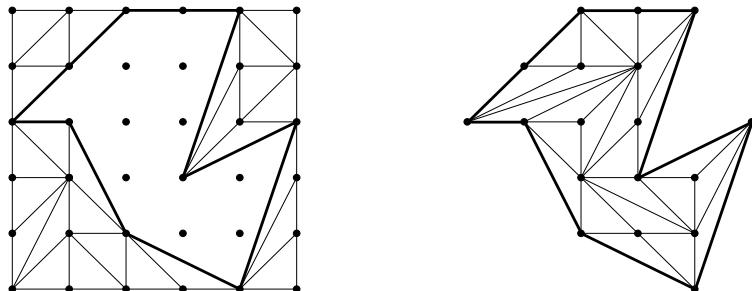


FIGURE 5. *Example of a triangulation of a lattice polygon embedded in a 5×5 square. Left: the polygon edges (in bold), along with the edges of an arbitrary triangulation outside the polygon, are fixed as constraints. Right: a triangulation consistent with the constraints.*

2.2. The flip graph. Any interior edge σ_x lies in two triangles. From area considerations, the third vertex of each of these triangles is an integer point on the line parallel to σ_x that passes through the closest integer point on either side of σ_x (see Figure 6(a)). The integer points on these lines occur periodically at intervals equal to the length of σ_x , and are positioned symmetrically on either side of σ_x , so that x is the midpoint of the line segment joining a point to its symmetric pair. The two triangles containing σ_x form a quadrilateral; this quadrilateral is convex iff the third vertices of the triangles are a symmetric pair. In this case σ_x is the diagonal of a parallelogram, and thus can be replaced by the other diagonal to yield another valid triangulation: σ_x is then said to be *flippable*. We observe that, for any σ_x that is not of minimum length, there is a unique parallelogram in which σ_x is the longer diagonal; we call this the *minimal parallelogram* of σ_x . Thus there is a unique flip that takes σ_x to a *shorter* edge. (There are many possible flips that make σ_x longer.) See Figure 6(b,c).

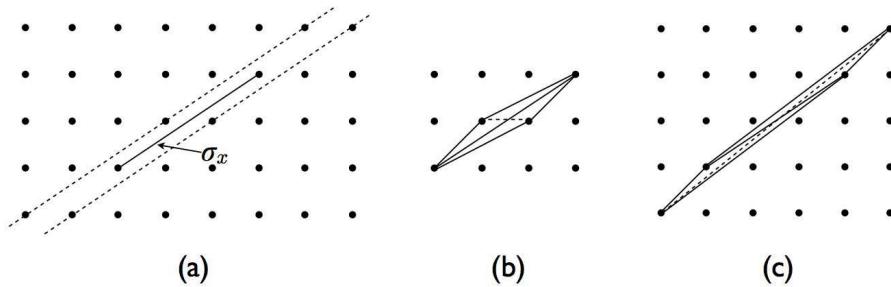


FIGURE 6. *Flips of an edge σ_x .* (a) Possible locations of the third vertex of the triangles containing σ_x are the integer points on the dotted lines. (b) Unique flip of σ_x to a shorter edge (σ_x is the diagonal of its minimal parallelogram). (c) A flip of σ_x to a longer edge.

We observe also that flipping an edge cannot change its slope from positive to negative (or vice versa), unless the edge is a unit diagonal. Thus any sequence of flips that changes the sign of the slope of an edge must pass through the minimum configuration of the edge (horizontal, vertical, or unit diagonal). Figure 7 shows how this is achieved.

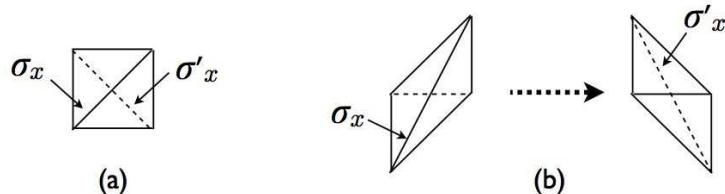


FIGURE 7. *Changing the sign of the slope of σ_x .* (a) If x is Type 2, the change can be effected only when σ_x is a unit diagonal. (b) If x is Type 1, σ_x must first become horizontal (or vertical), and then be flipped after the orientation of its surrounding parallelogram has been changed. In both cases the new configuration is denoted σ'_x .

The flipping operation induces a graph on vertex set Ω known as the “flip graph”, in which two triangulations are adjacent iff they differ by one edge flip. The flip graph is

well known to be connected. To see this, note that a longest edge in any non-ground-state triangulation is always flippable to a shorter edge, because if σ_x is the longest edge in both of its triangles then the triangles must form the minimal parallelogram of σ_x . Hence, for any initial triangulation σ , there exists a sequence of flips that reaches a ground state. Since flips are reversible, any triangulation is also reachable from a ground state. And finally, any ground state is reachable from any other by flipping unit diagonals.

For triangulations with constraints we get a flip graph on the smaller set $\Omega(\eta, \Delta)$. As we shall see in Section 3.2, in the presence of constraints there is still a well-defined ground state in which all configurations σ_x are of minimal length (subject to the constraints). Moreover, it can be shown (see Lemma 3.7 below) that any longest edge that is not in its ground state configuration is flippable to a shorter edge, and hence the flip graph remains connected in this case. In Proposition 3.8 we shall see that the flip graph actually enjoys a very strong structural property, which allows one to exactly compute the shortest path distance (or *flip distance*) between any pair of triangulations.

2.3. Weighted triangulations and Glauber dynamics. We associate with a triangulation σ a *weight* $\lambda^{|\sigma|}$, where λ is a positive real parameter and $|\sigma|$ denotes the *total length* of σ , i.e., $|\sigma| = \sum_x |\sigma_x|$; here $|\cdot|$ denotes the ℓ_1 (Manhattan) metric. These weights induce a probability distribution μ on Ω (or, in the presence of constraints, on $\Omega(\eta, \Delta)$) via

$$\mu(\sigma) = \frac{\lambda^{|\sigma|}}{Z}, \quad (2.1)$$

where the normalizing factor $Z = \sum_\sigma \lambda^{|\sigma|}$ is the *partition function*. We refer to (2.1) as the *Gibbs distribution*. The case $\lambda = 1$ corresponds to the uniform distribution on triangulations; the cases $\lambda < 1$ (respectively, $\lambda > 1$) favor shorter (respectively, longer) triangulations. Note that when $\lambda < 1$ the triangulations of maximum weight are precisely the *ground state* triangulations, i.e., those that minimize $|\sigma|$.²

The flip graph forms the basis of a local Markov chain (or *Glauber dynamics*) on Ω (or, more generally, on $\Omega(\eta, \Delta)$) as follows. In state $\sigma \in \Omega$, pick a point $x \in \Lambda$ uniformly at random³; if the edge σ_x is flippable to edge σ'_x (producing a new triangulation σ'), then flip it with probability

$$\frac{\mu(\sigma')}{\mu(\sigma') + \mu(\sigma)} = \frac{\lambda^{|\sigma'_x|}}{\lambda^{|\sigma'_x|} + \lambda^{|\sigma_x|}},$$

else do nothing. Since the flip graph is connected, this so-called “heat-bath” dynamics defines an ergodic Markov chain on Ω (or on $\Omega(\eta, \Delta)$) that is reversible with respect to μ . Hence the dynamics converges to the stationary distribution μ . We will analyze convergence to stationarity via the standard notion of *mixing time*, defined by

$$T_{\text{mix}} = \inf \left\{ t \in \mathbb{N} : \max_{\sigma \in \Omega} \|p^t(\sigma, \cdot) - \mu\| \leq 1/4 \right\}, \quad (2.2)$$

where $p^t(\sigma, \cdot)$ denotes the distribution after t steps when the initial state is σ , and $\|\nu - \mu\| = \frac{1}{2} \sum_{\sigma \in \Omega} |\nu(\sigma) - \mu(\sigma)|$ is the usual total variation distance between two distributions μ, ν .

²This explains our choice of the term “ground state”; in statistical physics, ground states are configurations of minimum energy, and thus maximum weight in the Gibbs distribution. The term “ground state” is less appropriate when $\lambda > 1$, but our focus in this paper is mainly on the case $\lambda < 1$.

³In the presence of constraints, we pick x u.a.r. from the midpoints of *non-constraint* edges.

3. STRUCTURAL PROPERTIES

In this section we establish some basic geometric properties of lattice triangulations and the flip graph. Throughout we will work with arbitrary constraints, so that our results apply in particular to triangulations of any lattice polygon.

3.1. The minimal parallelogram and excluded region. We begin with a useful property of the lattice \mathbb{Z}^2 .

Proposition 3.1. *Let a, b be positive and coprime, and consider the infinite line through $(0, 0)$ of slope (a, b) . Then the points of \mathbb{Z}^2 that are closest to this line (and are not on the line) are at horizontal distance a^{-1} and vertical distance b^{-1} on either side of the line.*

Proof. Consider two successive integer points on the line (say, $(0, 0)$ and (a, b)), and the a horizontal lines at integer heights $0, 1, 2, \dots, a - 1$. On each such horizontal line, consider the closest integer point to the line that lies strictly to the right of the line. (By symmetry, the points on the left of the line behave similarly.) It is readily verified that the horizontal distances of these points from the line are all distinct elements of the set $\{\frac{i}{a} : i = 1, 2, \dots, a\}$. Hence the closest of these points is at horizontal distance a^{-1} . This pattern repeats with period a throughout \mathbb{Z}^2 . The same argument applies to the vertical distances, with period b . Finally, note that a point at horizontal distance a^{-1} is also at vertical distance b^{-1} . \square

An immediate corollary is the following.

Proposition 3.2. *Let a, b be positive and coprime. The infinite family of parallel lines with slope (a, b) and horizontal separation a^{-1} (and thus vertical separation b^{-1}) partitions \mathbb{Z}^2 , in the sense that every integer point lies on one of the lines.*

Given any triangulation σ , consider an arbitrary edge σ_x that is not horizontal or vertical. Let its slope be (a, b) with a, b coprime, and denote its upper and lower endpoints by p_1, p_2 respectively. By Proposition 3.1 and its proof we may identify the unique closest integer point to the right of σ_x that lies horizontally and vertically between p_1 and p_2 . Call this point p_3 . By symmetry there is a corresponding closest point p_4 on the other side of σ_x , so that x is the midpoint of the line segment p_3p_4 . The points p_1, p_3, p_2, p_4 form the vertices of the *minimal parallelogram* of σ_x mentioned earlier: σ_x can be flipped to the shorter edge p_3p_4 iff all edges of the minimal parallelogram are present in the triangulation.

The following statement identifies a region centered on the minimal parallelogram that is free of integer points. (See Figure 8.)

Proposition 3.3. *Let σ_x be any non-horizontal and non-vertical triangulation edge with endpoints p_1, p_2 and minimal parallelogram $p_1p_3p_2p_4$. Then the shaded region in Figure 8 (the “excluded region”) does not contain any integer points in its interior.*

Note: The excluded region is the union of two strips with parallel sides that extend to the boundaries of Λ^0 . The intersection of the two strips is the minimal parallelogram of σ_x . The complement of the excluded region consists of four disjoint closed cones, which we shall call R_1, R_2, R_3, R_4 , as shown in Figure 8 (so the apex of R_i is p_i).

Proof. It suffices to show that each of the two pairs of parallel lines that bound the strips are adjacent lines in a family as in Proposition 3.2. Consider first the pair of lines through p_1p_4 and p_2p_3 . We need to verify that p_1 is a closest point to the line through p_2p_3 . W.l.o.g., assume that p_2 is the point $(0, 0)$ and p_1 is the point (a, b) , with a, b positive

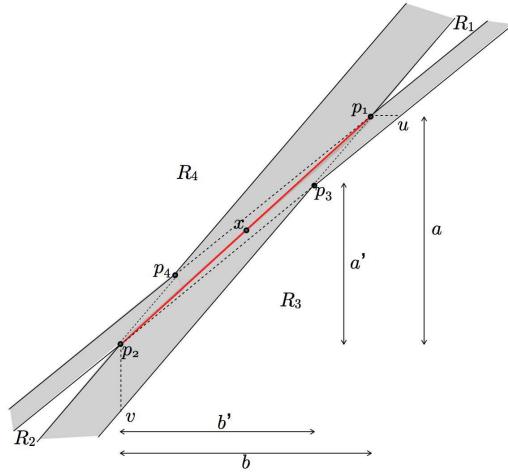


FIGURE 8. *The minimal parallelogram and excluded region of an edge σ_x .*

and coprime. Suppose p_3 is the point (a', b') , and recall that the horizontal distance of p_3 from σ_x is a^{-1} . By similarity of triangles, the horizontal distance of p_1 from the line through p_2p_3 (distance p_1u in Figure 8) is $\frac{a}{a'} \times \frac{1}{a} = \frac{1}{a'}$. Hence, by Proposition 3.1, p_1 is a closest point to this line, and by Proposition 3.2 this parallel strip contains no integer points.

An analogous argument using vertical separation shows that the distance p_2v is $\frac{1}{b'}$ and hence the other parallel strip, bounded by lines through p_2p_4 and p_1p_3 , also contains no integer points. This completes the proof. \square

3.2. The ground state lemma. We turn now to the structure of the *ground state* triangulations, i.e., those of minimum total length. These are the triangulations of maximum weight in the probability distribution (2.1) when $\lambda < 1$, and they play a central role in our analysis of spatial mixing in Section 4. In the absence of constraints, the ground state triangulations are very simple: every edge is either horizontal or vertical or a unit diagonal, so in particular the ground state is unique up to flipping of the unit diagonals. The presence of constraint edges, however, may change the ground state considerably. In this subsection we prove that the ground state remains (essentially) unique, and can be described easily (see Figure 9 for an example).

Lemma 3.4. [Ground State Lemma] *Given any set of constraint edges, the ground state triangulation is unique (up to possible flipping of unit diagonals), and can be constructed by placing each edge in its minimal length configuration consistent with the constraints, independent of the other edges.*

Lemma 3.4 will follow from a short sequence of observations. Given a set of constraints (η, Δ) , we call a midpoint x *constrained* if the minimal configuration of the edge at x consistent with the constraints is not horizontal or vertical or unit diagonal; or if the minimal configuration is a unit diagonal and its other unit diagonal configuration is not consistent.

Let σ_x be a consistent configuration for the edge at x with endpoints p_1, p_2 as in Figure 8. We say that σ_x is *spanned* by a constraint edge C if C crosses the line segment p_3p_4 in Figure 8 (i.e., the shorter diagonal of the minimal parallelogram of σ_x). Note that this necessarily implies that the endpoints of C lie in the regions R_1 and R_2 .

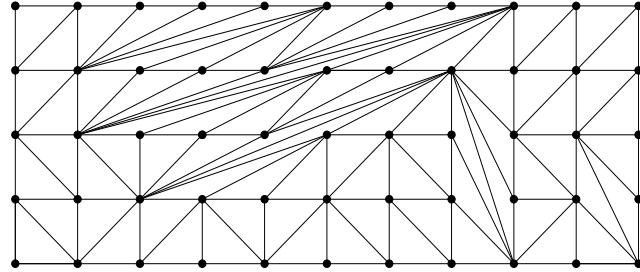


FIGURE 9. *Example of a ground state triangulation in the presence of the constraints given in Figure 4.*

Proposition 3.5. *Given a set of constraints, a consistent configuration σ_x for the edge at a constrained x is minimal if and only if σ_x is spanned by a constraint edge. Moreover, the minimal configuration $\sigma_x = \bar{\sigma}_x$ is unique.*

Proof. Assume first that σ_x is minimal and x is constrained. Then in particular the edge p_3p_4 must not be a consistent configuration for the edge at x (since either it would be a shorter configuration, or in the case of a unit diagonal it is explicitly prohibited). Hence a constraint edge must pass through p_3p_4 .

For the other direction, suppose σ_x is spanned by a constraint edge. We claim that any consistent configuration for the edge at x must have its endpoints in the cones R_1, R_2 defined by σ_x (see Figure 8). This immediately implies that σ_x is the unique minimal configuration for this edge. To see the above claim, let C be the spanning constraint edge. Since any configuration σ'_x for the edge at x must be obtained from σ_x by rotating about its midpoint x , and there are no integer points in the excluded region, the endpoints of σ'_x must lie either in R_1, R_2 or in R_3, R_4 . But an edge between R_3 and R_4 would necessarily cross C , so we are left only with the possibility R_1, R_2 . \square

Proposition 3.6. *Given a set of constraints, let $\bar{\sigma}_x, \bar{\sigma}_y$ be minimal configurations of any two distinct edges. Then $\bar{\sigma}_x, \bar{\sigma}_y$ do not cross.*

Proof. By Proposition 3.5, both $\bar{\sigma}_x$ and $\bar{\sigma}_y$ are either spanned by a constraint edge or the corresponding midpoints are unconstrained. If both x, y are unconstrained then clearly they do not cross. So assume w.l.o.g. that $\bar{\sigma}_x$ is spanned by some constraint edge C . Suppose for contradiction that $\bar{\sigma}_y$ crosses $\bar{\sigma}_x$. Then $\bar{\sigma}_y$ must have endpoints either in R_1, R_2 or in R_3, R_4 , where the R_i are the regions defined by $\bar{\sigma}_x$ as in Figure 8. As in the proof of Proposition 3.5, an edge between R_3 and R_4 would necessarily cross C , so $\bar{\sigma}_y$ must have endpoints in R_1, R_2 . But clearly this implies that $|\bar{\sigma}_y| > |\bar{\sigma}_x|$.

Now if y is constrained, then switching the roles of x and y in the above argument yields $|\bar{\sigma}_x| > |\bar{\sigma}_y|$, a contradiction. And if y is unconstrained then $\bar{\sigma}_y$ is horizontal, vertical or a unit diagonal, in which case it cannot be strictly longer than the constrained edge $\bar{\sigma}_x$, again a contradiction. \square

The Ground State Lemma (Lemma 3.4) now follows immediately from Propositions 3.5 and 3.6. Given any set of constraints (η, Δ) , we may thus speak of a minimal length triangulation $\bar{\sigma} = \{\bar{\sigma}_x\}_{x \in \Lambda}$ which is unique except for possible flipping of unit diagonals. We call $\bar{\sigma}$ a *ground state triangulation*.

3.3. Flip distance. We start with a proof of the claim in Section 2.2 that the flip graph remains connected in the presence of arbitrary constraints. (Recall that connectedness is

easily verified in the absence of constraints.) We shall also see that the diameter is always small, regardless of the constraints.

Lemma 3.7. *For any set of constraints (η, Δ) , the flip graph on triangulations in $\Omega(\eta, \Delta)$ is connected. Moreover, its diameter is $O(mn(m + n))$.*

Proof. By the Ground State Lemma (Lemma 3.4), there is a well-defined family of ground state triangulations for any such (η, Δ) , which differ only up to flipping of unit diagonals. As in the unconstrained case, it suffices to show that a ground state is reachable from any triangulation in $\Omega(\eta, \Delta)$, since by reversibility this implies that any triangulation is reachable from a ground state, and ground states are obviously reachable from each other. Since in any ground state triangulation all edges have minimal possible length, it suffices in turn to prove that, in any non-ground-state triangulation, there is an edge that is flippable to a shorter edge. This we now do.

Let $\sigma \in \Omega(\eta, \Delta)$ be an arbitrary non-ground-state triangulation, and let σ_x be some edge of σ that is not in ground state. If σ_x is flippable to a shorter edge we are done; if not, then σ_x cannot be the longest edge in both of its triangles, so we can find a *longer* edge, σ_y , that shares an endpoint with σ_x . The crucial observation (see the next paragraph) is that σ_y must also not be in ground state. Thus we can iterate this process, finding a connected sequence of edges of increasing length that are not in ground state, until we find one that is longest in both its triangles. This edge must be flippable to a shorter edge and we are done.

To conclude, we verify the above claim that σ_y must not be in ground state. This follows from the fact that, in any triangle, if the longest edge is in ground state then so are the other two edges. To see this, simply note that the longest edge in a triangle spans the other two edges; this can be seen because the triangle contains no lattice points, so the longest edge must pass between each of the other edges and its closest lattice point on the same side of the edge (see Figure 8).

Finally, the assertion concerning the diameter follows immediately from the fact that the maximum length of a triangulation is $O(mn(m + n))$, which is clearly an upper bound on the number of edge-shortening flips needed to reach a ground state triangulation. \square

Next, we shall compute the *flip distance* between two arbitrary triangulations in the flip graph, i.e., the minimal number of flips required to obtain one triangulation from the other. As usual, fix a set of constraints (η, Δ) . Given a midpoint x , let Ω_x denote the set of all possible values of the edge at x that are consistent with the constraints. This set contains the ground state $\bar{\sigma}_x$ (or, in the case where x is unconstrained and of Type 2, the pair of ground states corresponding to two opposite unit diagonals). Say that two elements $\sigma_x, \sigma'_x \in \Omega_x$ are *neighbors*, written $\sigma_x \sim \sigma'_x$, if σ_x is flippable to σ'_x within a valid triangulation $\sigma \in \Omega(\eta, \Delta)$. Recall that for any $\sigma_x \in \Omega_x$ there is at most one $\sigma'_x \sim \sigma_x$ such that $|\sigma_x| \geq |\sigma'_x|$. If σ_x is not in ground state, then an edge $\sigma'_x \in \Omega_x$ with $\sigma'_x \sim \sigma_x$ and $|\sigma_x| > |\sigma'_x|$ necessarily exists. This follows from the process described in the proof of Lemma 3.7, since σ_x must eventually be flippable to a shorter edge. These observations prove that the resulting graph with vertex set Ω_x is a tree rooted at the ground state $\bar{\sigma}_x$ (or, when the ground state is not unique, at a pair of neighboring ground states corresponding to unit diagonals). Given $\sigma_x, \tau_x \in \Omega_x$, let $\kappa(\sigma_x, \tau_x)$ be the distance between σ_x, τ_x in the tree, i.e., the minimal number of flips of the edge at x required to change its configuration from σ_x to τ_x (regardless of the disposition of the other edges). Notice that any path between σ and τ in the flip graph must contain all of the above flips for every x , and thus have length at least $\sum_x \kappa(\sigma_x, \tau_x)$. We call these the *indispensable* flips for the pair σ, τ .

The next result expresses a fundamental structural property of the flip graph: there is a path between any two triangulations that consists only of indispensable flips. Note that this path always exists, independent of the constraints. Although we shall not require this result in our subsequent analysis, we include it here for completeness.

Proposition 3.8. *The flip distance between any two triangulations $\sigma, \tau \in \Omega(\eta, \Delta)$ is equal to $\sum_x \kappa(\sigma_x, \tau_x)$.*

Proof. Given any pair of triangulations $\sigma, \tau \in \Omega(\eta, \Delta)$ with $\sigma \neq \tau$, we show how to perform an indispensable flip in either σ or τ to produce a new pair σ', τ' . Since the set of indispensable flips for σ', τ' differs from that for σ, τ only in the single flip just performed, iteration of this procedure until $\sigma = \tau$ produces a path of indispensable flips between σ and τ .

To justify the above claim, let $\Delta' = \{x \in \Lambda : \sigma_x = \tau_x\}$; note that Δ' contains Δ , and $(\Delta, \eta) = (\Delta, \sigma)$. We will view (Δ', σ) as an expanded set of constraints. Let x_* denote a point $x \in \Lambda \setminus \Delta'$ such that

$$|\sigma_{x_*}| = \max\{|\sigma_x| : x \in \Lambda \setminus \Delta'\}.$$

Similarly, let y_* denote a maximal edge in τ outside Δ' . Assume w.l.o.g. that $|\sigma_{x_*}| \geq |\tau_{y_*}|$ (else the same argument follows with the roles of σ_{x_*} and τ_{y_*} interchanged). We claim that an indispensable flip can be performed on σ_{x_*} .

Note first that $|\sigma_{x_*}| \geq |\tau_{y_*}| \geq |\tau_{x_*}|$, so since ground states are minimal and $\sigma_{x_*} \neq \tau_{x_*}$, the only way σ_{x_*} can be in ground state is if $|\sigma_{x_*}| = |\tau_{x_*}|$ and σ_{x_*}, τ_{x_*} are opposite unit diagonals; in this case an indispensable flip of σ_{x_*} must be possible (since it could be prevented only by a longer non-constraint edge, and σ_{x_*} is assumed to be maximal) and takes us directly to τ_{x_*} . Otherwise, σ_{x_*} is not in ground state, so from the proof of Lemma 3.7 since σ_{x_*} is a longest such edge it is flippable to a shorter edge. To see that this flip is indispensable, note that $\sigma_{x_*} \neq \tau_{x_*}$ and $|\sigma_{x_*}| \geq |\tau_{x_*}|$, so the path from σ_{x_*} to τ_{x_*} in the tree Ω_x must include the shortening flip at σ_{x_*} . This completes the proof. \square

3.4. The influence region. The following notion of “influence region” of an edge will play a key role in our analysis of decay of correlations under the Gibbs distribution for small λ . As usual we fix an arbitrary set of constraints (η, Δ) . Let $\bar{\sigma}$ be a ground state triangulation compatible with the constraints as in Lemma 3.4. Suppose we change the configuration of the edge at some point x to a non-ground-state value $\sigma_x \neq \bar{\sigma}_x$ (still consistent with the constraints). We would like to identify the minimal region “affected by” the new edge σ_x . We call this region the *influence region* of σ_x , and denote it $\Upsilon(\sigma_x)$. In this section we give two definitions of $\Upsilon(\sigma_x)$ —one implicit, the other constructive—and prove their equivalence.

For the first definition, we define $\Upsilon(\sigma_x)$ as the smallest region $S \subseteq \mathbb{R}^2$ made up of ground state triangles such that σ_x is contained in S . (When $\sigma_x = \bar{\sigma}_x$, we define $\Upsilon(\sigma_x)$ to be empty. Note that the ground state triangles in S may differ from those in $\bar{\sigma}$ up to flips of unit diagonals.) Observe that, since the boundary of $\Upsilon(\sigma_x)$ must consist of ground state edges that do not cross σ_x , the presence of σ_x has no effect on the ground state outside this region.

Now we describe an equivalent construction of the set $\Upsilon(\sigma_x)$ (refer to Figure 10). In the setting in the first paragraph above, let $\bar{\sigma}(\sigma_x)$ denote the new ground state triangulation when the edge σ_x is added as an additional constraint. Consider the region $\Upsilon^*(\sigma_x) \subseteq \mathbb{R}^2$ defined by successive addition of triangles as follows. Let y_1, y_2 and z_1, z_2 denote the midpoints of the edges of the minimal parallelogram of σ_x , so that this parallelogram is

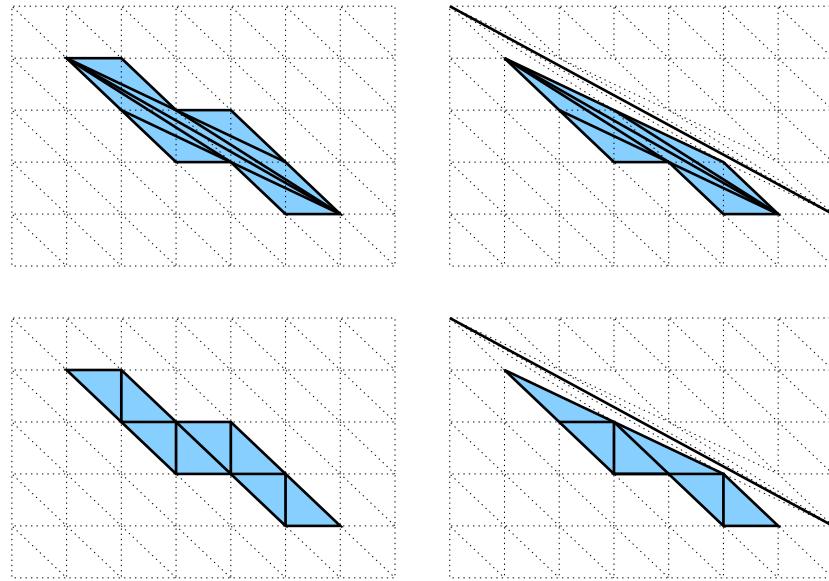


FIGURE 10. The influence region $\Upsilon(\sigma_x)$ when σ_x is an edge of slope $(-3, 5)$ (with midpoint in the center of the box). Left: in the case of no constraints. Right: in the presence of a constraint edge of slope $(-4, 7)$. Above: the triangles inside correspond to the second definition of the influence region. Below: the triangles inside correspond to the first definition of the influence region. In the background, a ground state triangulation (dotted lines).

made up of the two triangles $T_1 = (\sigma_x, \bar{\sigma}_{y_1}(\sigma_x), \bar{\sigma}_{y_2}(\sigma_x))$ and $T_2 = (\sigma_x, \bar{\sigma}_{z_1}(\sigma_x), \bar{\sigma}_{z_2}(\sigma_x))$. Note that these triangles are consistent with the constraints: indeed, if one of the constraint edges C were to cross an edge of one of these triangles, then C would span σ_x ; and by Proposition 3.5 this would imply that σ_x is minimal, contradicting the assumption that $\sigma_x \neq \bar{\sigma}_x$. Thus, we can add the triangles T_1 and T_2 . Next, we add a further triangle for each edge $\bar{\sigma}_{y_i}(\sigma_x)$, $i = 1, 2$, such that $\bar{\sigma}_{y_i}(\sigma_x) \neq \bar{\sigma}_{y_i}$ and for each edge $\bar{\sigma}_{z_i}(\sigma_x)$, $i = 1, 2$, such that $\bar{\sigma}_{z_i}(\sigma_x) \neq \bar{\sigma}_{z_i}$. For edges that are already in their original ground state we add no further triangle. More precisely, if $\bar{\sigma}_{y_1}(\sigma_x) \neq \bar{\sigma}_{y_1}$, then define midpoints y_{11}, y_{12} and the triangle $T_{11} = (\bar{\sigma}_{y_1}(\sigma_x), \bar{\sigma}_{y_{11}}(\sigma_x), \bar{\sigma}_{y_{12}}(\sigma_x))$, which is one half of the minimal parallelogram of $\bar{\sigma}_{y_1}(\sigma_x)$. The same reasoning as above shows that this triangle is consistent with the constraints as long as $\bar{\sigma}_{y_1}(\sigma_x) \neq \bar{\sigma}_{y_1}$. Similarly, define y_{21}, y_{22} and $T_{12} = (\bar{\sigma}_{y_2}(\sigma_x), \bar{\sigma}_{y_{21}}(\sigma_x), \bar{\sigma}_{y_{22}}(\sigma_x))$, which is one half of the minimal parallelogram of $\bar{\sigma}_{y_2}(\sigma_x)$. The same construction can be applied to the other side of σ_x to define the points $z_{11}, z_{12}, z_{21}, z_{22}$ and the associated triangles T_{21}, T_{22} . This branching procedure is iterated until one of the midpoints, say w , is such that $\bar{\sigma}_w(\sigma_x) = \bar{\sigma}_w$. In this case that branch is stopped and we continue with the other available branches, if any. $\Upsilon^*(\sigma_x)$ is defined as the union of all triangles added in this way.

Notice that all triangles added contain at least one edge that is not in ground state. However, the boundary of $\Upsilon^*(\sigma_x)$ is made up only of ground state edges. Therefore, $\Upsilon^*(\sigma_x)$ is also a union of ground state triangles. (See Figure 10.)

Lemma 3.9. *For any point x and any value of the edge σ_x , $\Upsilon^*(\sigma_x) = \Upsilon(\sigma_x)$.*

Proof. We can assume $\sigma_x \neq \bar{\sigma}_x$, since otherwise $\Upsilon^*(\sigma_x) = \Upsilon(\sigma_x) = \emptyset$. Since $\Upsilon^*(\sigma_x)$ is a union of ground state triangles that contains σ_x , we have $\Upsilon(\sigma_x) \subseteq \Upsilon^*(\sigma_x)$. Next, suppose

for contradiction that $\Upsilon(\sigma_x)$ is strictly contained in $\Upsilon^*(\sigma_x)$. Then there must be a point z such that $\bar{\sigma}_z$ is a boundary edge of $\Upsilon(\sigma_x)$ but z is an interior point of $\Upsilon^*(\sigma_x)$. This implies that z can achieve its ground state value $\bar{\sigma}_z$ even in the presence of σ_x , but the fact that z is an interior point of $\Upsilon^*(\sigma_x)$ implies that $\bar{\sigma}_z(\sigma_x) \neq \bar{\sigma}_z$ (otherwise the branching procedure defining $\Upsilon^*(\sigma_x)$ would have stopped at z and z would not be an interior point). Thus, such a z cannot exist. \square

4. SPATIAL MIXING FOR SMALL λ

In this section we prove some fundamental properties of the Gibbs distribution μ for sufficiently small values $\lambda < 1$. Since long edges are penalized in this regime, one might expect that the probability that an edge is longer than its ground state value should decrease exponentially in this excess length. We prove this property in Corollary 4.2 below, as a consequence of a rather more general exponential tail property (Lemma 4.1). We then go on to establish exponential decay of correlations (or *spatial mixing*) in the same regime (see Theorem 4.4), as advertised in Theorem 1.2 in the Introduction. It is worth emphasizing that all these properties are shown to hold uniformly in the constraints (η, Δ) , and thus in particular for triangulations of arbitrary lattice polygons.

We begin with some additional terminology. Fix a set of constraint edges. For any triangulation σ consistent with the constraints, we call a triangle in σ a *G-triangle* if all of its edges are in ground state (with respect to the constraints), and a *B-triangle* otherwise. Two triangles are *connected* if they share an edge, and **-connected* if they share a vertex. Note that every maximal connected region of B-triangles is bounded by ground state edges (see Figure 11). For any midpoint x , we denote by $S_x(\sigma)$ the maximal connected

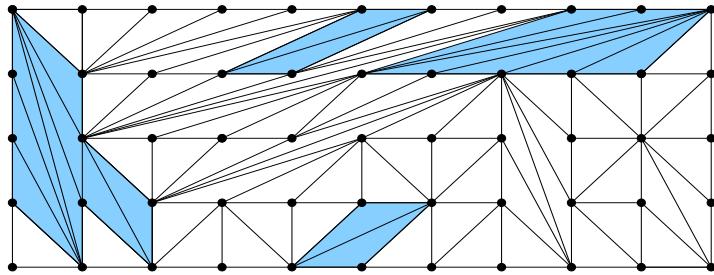


FIGURE 11. *The triangulation with constraints of Figure 4, showing its four connected regions of B-triangles (shaded). All unshaded triangles are G-triangles. Compare also with Figure 9.*

component of B-triangles containing the edge at x (so $S_x(\sigma) = \emptyset$ iff σ_x is contained in two G-triangles). Notice that $S_x(\sigma)$ need not be simply connected (i.e., there may be holes made up of G-triangles), and the boundary of $S_x(\sigma)$ consists of ground state edges. We call $\mathcal{S} = \mathcal{S}(\sigma)$ the union of all B-triangles in σ , i.e., $\mathcal{S}(\sigma) = \bigcup_x S_x(\sigma)$.

For any region $S \subseteq \Lambda^0$ that is the union of ground state triangles, let $\Lambda(S)$ denote the set of interior midpoints of S , namely the collection of $z \in \Lambda \cap S$ that are not on the boundary of S . Notice that any triangulation of the region S is obtained by specifying the edges σ_z , $z \in \Lambda(S)$. We define

$$\Phi(S, \sigma) := \sum_{z \in \Lambda(S)} |\sigma_z|.$$

In particular, $\Phi(S_x(\sigma), \sigma)$ is just the total length of the triangulation inside the region $S_x(\sigma)$. Clearly, for any σ , we have $\Phi(S_x(\sigma), \sigma) \geq \Phi(S_x(\sigma), \bar{\sigma})$, where $\bar{\sigma}$ denotes a ground state triangulation. The inequality is strict whenever $S_x(\sigma) \neq \emptyset$. It will be useful to consider the quantity

$$\phi_x(\sigma) := \Phi(S_x(\sigma), \sigma) - \Phi(S_x(\sigma), \bar{\sigma}). \quad (4.3)$$

The next lemma shows that, for sufficiently small λ , ϕ_x has exponentially decaying tail probability.

Lemma 4.1. *There exists $\lambda_0 \in (0, 1)$ such that, for all sets of constraints, all points $x \in \Lambda$, all $\lambda > 0$ and all $k \in \mathbb{N}$,*

$$\mu(\phi_x(\sigma) = k) \leq (\lambda/\lambda_0)^k. \quad (4.4)$$

Proof. We have

$$\mu(\phi_x(\sigma) = k) = \sum_S \mu(\phi_x(\sigma) = k; S_x(\sigma) = S).$$

Let Z denote the partition function, i.e., $Z = \sum_{\sigma \in \Omega} \lambda^{|\sigma|}$. For a fixed region S that is the union of ground state triangles, write $\tau \sim S$ if τ is a triangulation of S , and let $Z_S = \sum_{\tau \sim S} \lambda^{|\tau|}$ be the restricted partition function within S . As a convention, for $\tau \sim S$, we do not include the boundary edges in τ , so that

$$|\tau| = \Phi(S, \tau) = \sum_{z \in \Lambda(S)} |\tau_z|.$$

Notice that for any region S as above we have the bound

$$Z \geq \lambda^{|\partial S|} Z_S Z_{S^c},$$

where $|\partial S|$ denotes the length of the boundary of S . For a triangulation $\sigma \sim S$ we also write $\sigma \sim^* S$ if every triangle in $\sigma \sim S$ has at least one edge not in ground state. Then,

$$\begin{aligned} \mu(\phi_x(\sigma) = k; S_x(\sigma) = S) &= \frac{\sum_{\sigma} \lambda^{|\sigma|} \mathbf{1}_{\phi_x(\sigma) = k; S_x(\sigma) = S}}{Z} \\ &\leq \frac{\lambda^{|\partial S|} Z_S Z_{S^c} \sum_{\sigma \sim S} \lambda^{\Phi(S, \sigma)} \mathbf{1}_{\Phi(S, \sigma) = \Phi(S, \bar{\sigma}) + k; \sigma \sim^* S}}{Z} \\ &\leq \frac{\sum_{\sigma \sim S} \lambda^{\Phi(S, \sigma)} \mathbf{1}_{\Phi(S, \sigma) = \Phi(S, \bar{\sigma}) + k; \sigma \sim^* S}}{Z_S}. \end{aligned}$$

Next, we can use the trivial bound $Z_S \geq \lambda^{\Phi(S, \bar{\sigma})}$, so that

$$\mu(\phi_x(\sigma) = k) \leq \sum_S \sum_{\sigma \sim S} \lambda^k \mathbf{1}_{\Phi(S, \sigma) = \Phi(S, \bar{\sigma}) + k; \sigma \sim^* S}.$$

Since $\sigma \sim^* S$, we have $\Phi(S, \sigma) - \Phi(S, \bar{\sigma}) \geq \frac{1}{2}|S|$, where $|S|$ stands for the number of triangles in S . (Recall that, since all triangles are unimodular, this number does not depend on the particular triangulation $\sigma \sim S$). Then, for any $\lambda_0 > 0$ we have

$$\mu(\phi_x(\sigma) = k) \leq (\lambda/\lambda_0)^k \sum_S \sum_{\sigma \sim S} \lambda_0^{|S|/2}. \quad (4.5)$$

From Anclin's bound [1] mentioned in Section 1.3, there are at most $2^{3|S|/2}$ triangulations of a region S , since the number of interior edges is at most $3|S|/2$. Moreover, the number of connected regions S containing the point x as above and with a given size $|S|$ is bounded by $9^{|S|}$. To see this, observe that every such S must be a collection of $|S|$ connected ground

state triangles, so it suffices to count the number of connected sets made up of ℓ triangles in a valid triangulation, which is at most $3^{2\ell}$ since each triangle is connected to at most 3 other triangles and a general connected set can be explored by a path that visits the same point at most twice. In conclusion, the claim follows by summing in (4.5):

$$\mu(\phi_x(\sigma) = k) \leq (\lambda/\lambda_0)^k \sum_S 2^{3|S|/2} \lambda_0^{|S|/2} \leq (\lambda/\lambda_0)^k \sum_{\ell=2}^{\infty} C_0^\ell \lambda_0^{\ell/2}, \quad (4.6)$$

where $C_0 = 9\sqrt{8}$ and we use the fact that the minimal non trivial S has size at least 2. It suffices to take $\lambda_0 > 0$ such that (e.g.) $\lambda_0^{1/2} C_0 \leq 1/2$ to ensure that the last summation is less than 1. \square

We are now able to prove the exponential tail bound advertised at the start of this section. Notice that if $\sigma \in \Omega$ is such that $|\sigma_x| = |\bar{\sigma}_x| + k$, for some point $x \in \Lambda$, then necessarily $\phi_x(\sigma) \geq k$. Hence from Lemma 4.1 we immediately have:

Corollary 4.2. *Let λ_0 be as in Lemma 4.1. For every point $x \in \Lambda$ and every $k \in \mathbb{N}$,*

$$\mu(|\sigma_x| = |\bar{\sigma}_x| + k) \leq (\lambda/\lambda_0)^k. \quad (4.7)$$

Recall that $\mathcal{S} = \mathcal{S}(\sigma)$ is the random subset of the plane defined as the union of all B-triangles in σ . The set \mathcal{S} naturally describes the deviation of σ from the ground state $\bar{\sigma}$. Our next result, based on a slight variation of the argument of Lemma 4.1, establishes that for small λ the probability that \mathcal{S} contains any particular region V decays exponentially in the size of V . Since the boundary of \mathcal{S} consists of ground state edges, \mathcal{S} can always be described as a collection of ground state triangles. Thus we naturally also specify V as a collection of ground state triangles. We write $|V|$ for the number of triangles in V .

Lemma 4.3. *Let λ_0 be as in Lemma 4.1. For every $V \subseteq \Lambda^0$ that is the union of ground state triangles, we have*

$$\mu(V \subseteq \mathcal{S}) \leq (4\lambda/\lambda_0)^{|V|/2}. \quad (4.8)$$

Proof. Suppose first that V is connected. Let T be a ground state triangle in V and let \mathcal{S}_0 denote the connected component of \mathcal{S} containing T . Then

$$\mu(V \subseteq \mathcal{S}) \leq \sum_{S: S \supseteq V} \mu(\mathcal{S}_0 = S),$$

where the sum is over all connected sets S of ground state triangles such that $S \supseteq V$. Reasoning as in the proof of Lemma 4.1 we have that

$$\mu(\mathcal{S}_0 = S) \leq (2^3 \lambda)^{|S|/2}. \quad (4.9)$$

Summing over all S as above and estimating by 9^ℓ the number of connected $S \ni T$ with $|S| = \ell$, with the notation of (4.6) we have

$$\mu(V \subseteq \mathcal{S}) \leq \sum_{\ell \geq |V|} C_0^\ell \lambda^{\ell/2} \leq (\lambda/\lambda_0)^{|V|/2}.$$

We turn to the case where V has several connected components, say V_1, \dots, V_n . Let T_1, \dots, T_n denote fixed ground state triangles such that $T_i \in V_i$, and let \mathcal{S}_i , $i = 1, \dots, n$ denote the random sets defined as the connected components of \mathcal{S} such that $\mathcal{S}_i \ni T_i$. Notice that the \mathcal{S}_i need not be distinct, since a single connected component of \mathcal{S} may contain more than one of the triangles T_i . However, the event $V \subseteq \mathcal{S}$ implies that there exist i_1, \dots, i_m , $1 \leq m \leq n$, and distinct connected components $\mathcal{S}_{i_1}, \dots, \mathcal{S}_{i_m}$ such that

$S_{i_k} \ni T_{i_k}$ and $\mathcal{S}_{i_k} = S_{i_k}$, for all $k = 1, \dots, m$. For any $\eta \in \Omega$ such that $\mathcal{S}_{i_k} = S_{i_k}$, $k = 1, \dots, m$, reasoning as in the proof of Lemma 4.1 and (4.9) we have

$$\mu(\mathcal{S}_{i_k} = S_{i_k}, \forall k = 1, \dots, m) \leq (2^3 \lambda)^{\sum_{k=1}^m |S_{i_k}|/2}.$$

Therefore, using $C_0 \lambda_0^{1/2} \leq 1/2$, $n \leq |V|$, and a union bound we arrive at

$$\begin{aligned} \mu(V \subseteq \mathcal{S}) &\leq \sum_{m=1}^n \binom{n}{m} \sum_{\ell_1 \geq 1} \dots \sum_{\ell_m \geq 1} 1_{\{\sum_{k=1}^m \ell_k \geq |V|\}} (C_0 \lambda^{1/2})^{\sum_{k=1}^m \ell_k} \\ &\leq \sum_{m=1}^n \binom{n}{m} (\lambda/\lambda_0)^{|V|/2} \leq 2^n (\lambda/\lambda_0)^{|V|/2} \leq (4\lambda/\lambda_0)^{|V|/2}. \end{aligned}$$

□

Finally, we use Lemma 4.3 to derive the spatial mixing estimate stated in Theorem 1.2 in the Introduction. This bound will be expressed in terms of a natural distance between the edge σ_x and any subset of midpoints $A \subseteq \Lambda$, which we now define.

Let σ_x be an edge consistent with the constraints (η, Δ) . Also, let $\bar{\sigma}(\sigma_x)$ denote the new ground state triangulation when edge σ_x is added to the constraints. Let μ^{σ_x} denote the Gibbs distribution on triangulations with this extra constraint, i.e., μ^{σ_x} is the probability μ conditioned on the event that the edge at x equals σ_x . Intuitively, the distributions μ^{σ_x} and μ should have very similar marginals at midpoints z that are far from the influence region $\Upsilon(\sigma_x)$, as defined in Section 3.4. To quantify this we introduce the following distance. Consider the graph $\mathcal{G}(\sigma_x)$ whose vertices are the triangles in $\bar{\sigma}(\sigma_x)$, where two triangles are neighbors iff they share an edge. Let $d(A, \sigma_x)$ denote the graph distance in $\mathcal{G}(\sigma_x)$ between the influence region $\Upsilon(\sigma_x)$ and the set of triangles spanned by edges $z \in A$ (i.e., the union of all triangles in $\bar{\sigma}(\sigma_x)$ containing the edges $\bar{\sigma}_z(\sigma_x)$, $z \in A$), where the distance between two sets of vertices of a graph is interpreted as the minimum distance between a pair of vertices, one from each set. In words, $d(A, \sigma_x)$ is the minimal number of triangles in $\bar{\sigma}(\sigma_x)$ needed to connect an edge in A to the influence region $\Upsilon(\sigma_x)$. Below we use $\|\mu^{\sigma_x} - \mu\|_A$ for the variation distance between the distributions μ^{σ_x} and μ at A , that is the variation distance between the marginals on edges $z \in A$. We write $|A|$ for the cardinality of A .

Theorem 4.4. *There exists $\lambda_1 \in (0, 1)$ such that, for all sets of constraints, all $x \in \Lambda$ and $A \subseteq \Lambda$ and all consistent values of the edge σ_x , we have*

$$\|\mu^{\sigma_x} - \mu\|_A \leq |A|(\lambda/\lambda_1)^{d(A, \sigma_x)}. \quad (4.10)$$

Proof. Observe that apart from the set $\Upsilon(\sigma_x)$ the ground states for μ^{σ_x} and μ coincide. When we talk about triangles below we always refer to the triangles in the common ground state triangulation (outside the region $\Upsilon(\sigma_x)$). Moreover, we identify A with the union of all triangles spanned by edges from A . Let \mathbb{P} denote the independent coupling of μ^{σ_x}, μ and let σ, σ' denote the associated random triangulations. Consider the event \mathcal{E} that there exists a $*$ -connected chain \mathcal{C} of ground state triangles that separates $\Upsilon(\sigma_x)$ from A such that both σ and σ' contain \mathcal{C} . From the Markov property of the Gibbs distributions μ^{σ_x}, μ it follows that they can be coupled at A in such a way that

$$\|\mu^{\sigma_x} - \mu\|_A \leq \mathbb{P}(\mathcal{E}^c). \quad (4.11)$$

Next, observe that the complementary event \mathcal{E}^c implies that there exists a connected chain \mathcal{D} of triangles connecting $\Upsilon(\sigma_x)$ and A such that, for every triangle $T \in \mathcal{D}$, either

$T \in \mathcal{S}(\sigma)$ or $T \in \mathcal{S}(\sigma')$, where $\mathcal{S}(\sigma)$ as usual denotes the set of B-triangles in a triangulation σ . Thus, given the set \mathcal{D} , either $\mathcal{S}(\sigma)$ or $\mathcal{S}(\sigma')$ contains at least half of the triangles in \mathcal{D} . By definition $|\mathcal{D}| \geq d(A, \sigma_x)$. In particular, from the union bound it follows that

$$\|\mu^{\sigma_x} - \mu\|_A \leq \sum_{\mathcal{D}} \sum_{V \subseteq \mathcal{D}: |V| \geq |\mathcal{D}|/2} [\mu^{\sigma_x}(V \subseteq \mathcal{S}) + \mu(V \subseteq \mathcal{S})]$$

where the first sum ranges over all connected subsets \mathcal{D} touching A and such that $|\mathcal{D}| \geq d(A, \sigma_x)$. From Lemma 4.3 it follows that for any set V , $\mu^{\sigma_x}(V \subseteq \mathcal{S}) \leq (4\lambda/\lambda_0)^{|V|/2}$ and $\mu(V \subseteq \mathcal{S}) \leq (4\lambda/\lambda_0)^{|V|/2}$, and therefore

$$\sum_{V \subseteq \mathcal{D}: |V| \geq |\mathcal{D}|/2} [\mu^{\sigma_x}(V \subseteq \mathcal{S}) + \mu(V \subseteq \mathcal{S})] \leq 2^{|\mathcal{D}|} (4\lambda/\lambda_0)^{|\mathcal{D}|/4}.$$

Using the fact that the number of connected \mathcal{D} touching a given triangle $T \subseteq A$ with $|\mathcal{D}| = k$ is at most 9^k , and bounding by $|A|$ the number of $T \subseteq A$, we get

$$\|\mu^{\sigma_x} - \mu\|_A \leq |A| \sum_{k \geq d(A, \sigma_x)} C_1^k (4\lambda/\lambda_0)^{k/4} \leq |A| (\lambda/\lambda_1)^{d(A, \sigma_x)/4},$$

for suitable constants $C_1, \lambda_1 > 0$. \square

5. UPPER BOUNDS ON THE MIXING TIME FOR $\lambda < 1$

We consider the Glauber dynamics on triangulations of the region $\Lambda^0 := \Lambda_{m,n}^0$, as defined in Section 2.3. We allow an arbitrary set of constraints (η, Δ) , so that our results apply in particular to triangulations of any lattice polygon. Our main result in this section is the following polynomial upper bound on the mixing time for all sufficiently small λ , which was stated as Theorem 1.3 in the Introduction. (Recall the definition of mixing time from equation (2.2).)

Theorem 5.1. *There exist constants $C > 0$ and $\lambda_0 \in (0, 1)$ such that, for all $\lambda \leq \lambda_0$, all $m, n \in \mathbb{N}$, and all constraints (η, Δ) ,*

$$T_{\text{mix}} \leq C m n (m + n). \quad (5.12)$$

We first establish a key coupling estimate and then turn to the proof of Theorem 5.1.

5.1. Path Coupling. Let σ and τ be two triangulations so that σ differs from τ only on the edge with midpoint x . Let $\alpha > 1$ be a fixed positive number. If σ_x and τ_x are opposite unit diagonals, define $\Delta(\sigma, \tau) = \alpha^2 - 1$; otherwise, let

$$\Delta(\sigma, \tau) = |\alpha^{|\sigma_x|} - \alpha^{|\tau_x|}|. \quad (5.13)$$

Note that Δ is defined only over pairs of triangulations that are adjacent in the flip graph. We extend Δ to all pairs of triangulations by defining $\Delta(\sigma, \tau)$ to be the shortest path distance between σ and τ in the flip graph, with edge lengths given by the above Δ values.

Lemma 5.2. *There exists $\alpha_0 > 1$ such that the following holds for all $\alpha \geq \alpha_0$, and $\lambda \alpha \leq 1$. Let σ, τ be two triangulations that differ in exactly one edge and let σ', τ' denote the random triangulation obtained after one step of the Markov chain. Then, there exists a coupling of σ' and τ' so that their expected distance satisfies*

$$\mathbb{E}[\Delta(\sigma', \tau')] \leq \Delta(\sigma, \tau) \left(1 - \frac{1}{2|\Lambda|}\right), \quad (5.14)$$

where Λ is the set of midpoints of non-constraint edges.

Proof. Let x be the midpoint so that σ_x and τ_x differ. Assume without loss of generality that $|\sigma_x| \geq |\tau_x|$. Since σ differs from τ by one flip at x , the edges σ_x and τ_x must be the two diagonals of the minimal parallelogram of σ_x . The coupling is such that the random midpoint $Y \in \Lambda$ chosen is the same for both triangulations. Thus, if $Y = x$ then $\sigma'_x = \tau'_x$ and $\Delta(\sigma', \tau') = 0$ with probability one. If instead $Y \neq x$, which happens with probability $(1 - 1/|\Lambda|)$, then $\Delta(\sigma', \tau') = \Delta(\sigma, \tau)$ unless Y is one of the edges of the minimal parallelogram of σ_x (since only these edges share a triangle with the differing edge at x); so it remains only to bound the expected change in distance induced by flips at the edges of the minimal parallelogram.

Assume first that $|\sigma_x| > |\tau_x|$. Note that in this case the length of σ_x can be written as $f_1 + f_2$, where f_1 and f_2 are the lengths of the other edges of the triangles containing σ_x in σ . Assume w.l.o.g. that $f_1 \geq f_2$; then we have $|\tau_x| = f_1 - f_2$. Let $y_1, y_2, z_1, z_2 \in \Lambda$ be the midpoints of the edges of the minimal parallelogram, so that $|\sigma_{y_1}| = |\sigma_{y_2}| = f_1 \geq f_2 = |\sigma_{z_1}| = |\sigma_{z_2}|$. We start with the case $Y = z_1$. Note that if σ_{z_1} (respectively, τ_{z_1}) is flippable, then τ_{z_1} (respectively, σ_{z_1}) is not flippable. Since $f_1 \geq f_2$, if σ_{z_1} is flippable it must flip to an edge σ'_{z_1} of length $|\sigma_x| + |\sigma_{y_1}| = 2f_1 + f_2$. Then, for $\alpha > 1$, $\lambda < 1$ with $\alpha\lambda \leq 1$, the expected increase in distance in this case is bounded by

$$\left(\alpha^{2f_1+f_2} - \alpha^{f_2}\right) \frac{\lambda^{2f_1+f_2}}{(\lambda^{2f_1+f_2} + \lambda^{f_2})} = \frac{\alpha^{f_2}(1 - \alpha^{-2f_1})(\lambda\alpha)^{2f_1}}{(\lambda^{2f_1} + 1)} \leq \alpha^{f_2}.$$

If instead τ_{z_1} is flippable, then it flips to an edge of length $|\tau_x| + |\tau_{y_1}| = 2f_1 - f_2$, resulting in an expected increase in distance of at most

$$\left(\alpha^{2f_1-f_2} - \alpha^{f_2}\right) \frac{\lambda^{2f_1-f_2}}{(\lambda^{2f_1-f_2} + \lambda^{f_2})} \leq \alpha^{f_2}(\lambda\alpha)^{2f_1-2f_2} \leq \alpha^{f_2}.$$

Now we turn to the case $Y = y_1$. Note that σ_{y_1} is either unflippable or flippable to a longer edge, while τ_{y_1} is either unflippable or flippable to a shorter edge. If σ_{y_1} is flippable, it flips to an edge of length $|\sigma_x| + |\sigma_{z_1}| = f_1 + 2f_2$, and the expected increase in distance is at most

$$\left(\alpha^{f_1+2f_2} - \alpha^{f_1}\right) \frac{\lambda^{f_1+2f_2}}{(\lambda^{f_1+2f_2} + \lambda^{f_1})} \leq \alpha^{f_1}(\lambda\alpha)^{2f_2} \leq \alpha^{f_1}.$$

If τ_{y_1} is flippable, it flips to an edge of length $|\tau_x| + |\tau_{z_1}| = |f_1 - 2f_2|$, which causes an expected increase in distance of at most

$$\left(\alpha^{|f_1-2f_2|} - \alpha^{f_1}\right) \leq \alpha^{|f_1-2f_2|} \leq \alpha^{f_1}.$$

Summing up all these contributions, we obtain

$$\mathbb{E}[\Delta(\sigma', \tau')] \leq \Delta(\sigma, \tau) \left(1 - \frac{1}{|\Lambda|}\right) + \frac{2\alpha^{f_2}}{|\Lambda|} + \frac{2\alpha^{f_1}}{|\Lambda|}.$$

This implies the desired bound in the lemma, since if α is large enough ($\alpha > 5$ suffices) we have

$$2\alpha^{f_1} + 2\alpha^{f_2} \leq \frac{1}{2}(\alpha^{f_1+f_2} - \alpha^{f_1-f_2}) = \frac{1}{2}\Delta(\sigma, \tau),$$

for any $f_1 > f_2 \geq 1$.

It remains to consider the case $|\sigma_x| = |\tau_x|$; note that here $\Delta(\sigma, \tau) = \alpha^2 - 1$. This case can only happen if σ_x, τ_x are opposite unit diagonals, so $|\sigma_x| = |\tau_x| = f_1 + f_2$ with

$f_1 = f_2 = 1$. Therefore, the total contribution from the edges of the minimal parallelogram (now a square) is bounded by

$$4(\alpha^3 - \alpha) \frac{\lambda^3}{(\lambda^3 + \lambda)} \leq 4\alpha^{-1}(\alpha^2 - 1)(\lambda\alpha)^2 \leq 4\alpha^{-1}\Delta(\sigma, \tau).$$

This gives the following bound on the expected distance after one flip:

$$\mathbb{E}[\Delta(\sigma', \tau')] \leq \Delta(\sigma, \tau) \left(1 - \frac{1}{|\Lambda|}\right) + \frac{4\alpha^{-1}\Delta(\sigma, \tau)}{|\Lambda|}.$$

The desired bound again follows by taking $\alpha \geq 8$. \square

Note: In the above proof it suffices to take $\alpha = 8$ and $\lambda \leq \alpha^{-1} = 1/8$. Hence Theorem 5.1 holds with $\lambda_0 = 1/8$.

5.2. Proof of Theorem 5.1. Once the estimate in Lemma 5.2 is available the argument is rather standard; see, e.g., [11, Theorem 14.6]. Indeed, from the triangle inequality and the definition of the metric $\Delta(\cdot, \cdot)$, it follows that the estimate (5.14) can be extended to *any* pair of triangulations σ, τ under a suitable coupling. This fact, together with the Markov property yields the bound

$$\mathbb{E}[\Delta(\sigma^{(k)}, \tau^{(k)})] \leq \Delta(\sigma, \tau) \left(1 - \frac{1}{2|\Lambda|}\right)^k \leq \Delta(\sigma, \tau) e^{-k/2|\Lambda|}, \quad (5.15)$$

where \mathbb{E} denotes expectation over the coupling of the random triangulations $\sigma^{(k)}, \tau^{(k)}$ obtained after k steps of the Markov chain started at σ, τ respectively. The diameter bound in Lemma 3.7, and the fact that the distance between any two adjacent triangulations is at most α^{m+n} , imply that $\Delta(\sigma, \tau) \leq e^{C_0(m+n)}$ for some constant $C_0 > 0$. Using also the fact that $\Delta(\sigma, \tau) \geq c$ if $\sigma \neq \tau$, for some constant $c = c(\alpha) > 0$, an application of Markov's inequality yields

$$\|p^k(\sigma, \cdot) - p^k(\tau, \cdot)\| \leq c^{-1} e^{C_0(m+n)-k/2|\Lambda|},$$

uniformly in the initial conditions σ, τ . Taking, e.g., $k = 4|\Lambda|C_0(m+n)$ concludes the proof. \square

Note: Observe from the above proof that Theorem 5.1 actually holds with the bound $T_{\text{mix}} \leq C|\Lambda|(m+n)$, where $|\Lambda|$ is the number of *non-constraint* edges. This bound may be much better in cases where there are many constraints.

5.3. The $1 \times n$ case for all $\lambda < 1$. Here we prove that, in the special case of the 1-dimensional region $\Lambda_{1,n}^0$, the bound of Theorem 5.1 holds for *all* $\lambda < 1$.

Theorem 5.3. *Set $m = 1$. For any $\lambda < 1$, there exists a constant $C > 0$ such that for all $n \in \mathbb{N}$, and for all constraints (η, Δ) ,*

$$T_{\text{mix}} \leq Cn^2. \quad (5.16)$$

To prove the theorem, we show that the path coupling argument of Section 5.1 works for all $\lambda < 1$ in this case. For 1D triangulations we may define the length of an edge as its horizontal length. Fix $\alpha > 1$. Let σ, τ be two triangulations that differ at exactly one edge x with $|\sigma_x| \geq |\tau_x|$. Then either the edges at x have equal length, in which case they are necessarily both unit diagonals, or they have lengths $\ell - 1$ (in τ) and $\ell + 1$ (in σ), for some integer $\ell \geq 1$. In the first case we set $\Delta(\sigma, \tau) = \alpha^2 - 1$, in the second case we set $\Delta(\sigma, \tau) = \alpha^{\ell+1}(1 - \alpha^{-2})$. As before, we extend Δ to the shortest-path metric on all pairs

in Ω . As detailed in Section 5.2, the estimate (5.16) will follow from the lemma below and the fact that the maximal distance between two arbitrary configurations is $e^{O(n)}$.

Lemma 5.4. *For any $\lambda < 1$ there exist constants $\alpha > 1$, $\delta > 0$ such that the following holds. Let σ, τ be two triangulations that differ at exactly one edge, and let σ', τ' denote the random triangulation obtained after one step. There exists a coupling of σ', τ' such that*

$$\mathbb{E} [\Delta(\sigma', \tau')] \leq (1 - \delta/n) \Delta(\sigma, \tau) \quad (5.17)$$

Proof. Say that the initial discrepancy is at midpoint x . In the coupling, we pick the same random midpoint X to be updated, so that $\mathbb{E}[d(\sigma', \tau')]$ is given by

$$\mathbb{E}[\Delta(\sigma', \tau')] = (1 - 1/n) \Delta(\sigma, \tau) + \frac{1}{n} U(\sigma, \tau),$$

where $U(\sigma, \tau)$ is the term coming from possible new discrepancies at $x \pm 1/2$. To compute the latter, we start with the case where $\sigma_x \neq \tau_x$ are both unit diagonals, so that $\Delta(\sigma, \tau) = \alpha^2 - 1$. In this case, the edges at $x \pm 1/2$ are necessarily vertical in both σ, τ and we can create a discrepancy at either $x \pm 1/2$ by letting the edge there increase to length 2. Only one of the two triangulations can flip at either $x \pm 1/2$. Thus, taking expectations, in this case we have

$$U(\sigma, \tau) \leq \frac{2\lambda^2}{1 + \lambda^2} (\alpha^2 - 1) = \frac{2\lambda^2}{1 + \lambda^2} \Delta(\sigma, \tau).$$

Next, consider the remaining case where $|\sigma_x| = \ell + 1$ and $|\tau_x| = \ell - 1$. Necessarily the edges at $x \pm 1/2$ have length ℓ in both σ, τ . Again, only one of the two triangulations can flip at either $x \pm 1/2$. In the worst case scenario, $x \pm 1/2$ can either both increase to $\ell + 2$, or both decrease to $\ell - 2$, or one can increase to $\ell + 2$ and the other decrease to $\ell - 2$. These produce respectively the terms

$$\begin{aligned} U(\sigma, \tau) &= \frac{2\lambda^2}{1 + \lambda^2} \alpha^{\ell+2} (1 - \alpha^{-2}) = \frac{2\alpha\lambda^2}{1 + \lambda^2} \Delta(\sigma, \tau); \\ U(\sigma, \tau) &= \frac{2}{1 + \lambda^2} \alpha^{\ell-1} (1 - \alpha^{-2}) = \frac{2}{\alpha(1 + \lambda^2)} \Delta(\sigma, \tau); \\ U(\sigma, \tau) &= \left(\frac{\alpha\lambda^2}{1 + \lambda^2} + \frac{1}{\alpha(1 + \lambda^2)} \right) \Delta(\sigma, \tau). \end{aligned}$$

Thus it suffices to show that, for any $\lambda < 1$, we can find $\alpha > 1$ and $\delta > 0$ such that

$$\max \left\{ \frac{2\alpha\lambda^2}{1 + \lambda^2}, \frac{2}{\alpha(1 + \lambda^2)} \right\} \leq 1 - \delta.$$

This is easily verified. \square

Note: The bound in Theorem 5.3 should be compared with the known bound $T_{\text{mix}} = \Theta(n^3 \log n)$ in the case of 1D triangulations with $\lambda = 1$ (and no constraints). This follows from a one-to-one correspondence between 1D triangulations and lattice paths, and the results of Wilson [23] on the mixing time of the flip dynamics for such paths.

6. LOWER BOUNDS ON THE MIXING TIME FOR $\lambda > 1$

As a general rule, our lower bounds on the mixing time are obtained by exploiting the fact that it takes a long time to change the orientation of an initially long edge. Below, we say that an edge of a triangulation is *positively oriented* (respectively, *negatively oriented*)

if its leftmost endpoint lies below (respectively, above) its rightmost endpoint. A vertical or horizontal edge is neither positively nor negatively oriented.

Our first result is an exponential lower bound on the mixing time (in the absence of constraints) that holds for any $\lambda > 1$. This proves Theorem 1.4 in the Introduction.

Theorem 6.1. *For every $\lambda > 1$, there exists $c > 0$ such that the mixing time of the Glauber dynamics on triangulations of $\Lambda_{m,n}^0$ in the absence of constraints satisfies $T_{\text{mix}} \geq \exp(c(m+n))$ for all $m,n \geq 1$.*

Proof. For any $A \subseteq \Omega$ we write

$$Z(A) = \sum_{\sigma \in A} \lambda^{|\sigma|}$$

for the partition function restricted to A . If ∂A denotes the set of $\sigma \in A$ that are adjacent to A^c in the flip graph, then the standard conductance bound (see, e.g., [11, Theorem 7.3]) will allow us to deduce the theorem once we find a set A such that $\mu(A) \leq 1/2$ and

$$\frac{Z(\partial A)}{Z(A)} \leq e^{-c(m+n)}. \quad (6.18)$$

Assume w.l.o.g. that the horizontal coordinate $n \geq m$ (else the same argument applies with the roles of n and m reversed). We define $A \subseteq \Omega$ as the set of triangulations $\sigma \in \Omega$ such that every internal midpoint with half-integer vertical coordinate v is not positively (resp., not negatively) oriented if $v + \frac{1}{2}$ is odd (resp., even). Notice that any $\sigma \in A$ has the ‘‘herringbone’’ structure illustrated in Figure 12. In particular, any triangulation in A consists of 1D triangulations in each horizontal layer. Moreover, if $\sigma \in A$ then all edges with integer vertical coordinate are horizontal and frozen (unflippable).

Now, the total length of any $\sigma \in A$ can be written as

$$|\sigma| = (m+1)n + (2n+1)m + \sum_{j=1}^m L(\sigma(j)),$$

where $(m+1)n$ is the total length of all horizontal (unit length) edges, $(2n+1)m$ is the sum of all vertical lengths of the edges of σ , and $L(\sigma(j))$ denotes the total horizontal length of the internal edges of the 1D triangulation $\sigma(j)$ in the j -th layer.

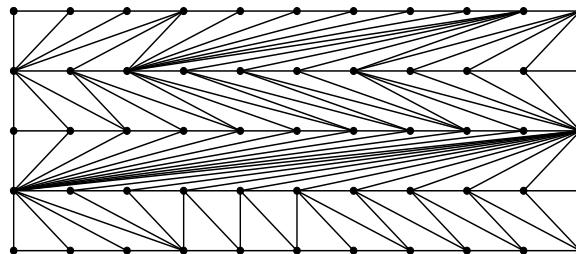


FIGURE 12. *Example of a triangulation of a 4×10 region, with the ‘‘herringbone’’ structure of the set A introduced in the proof of Theorem 6.1*

Below, we let $\Omega_{1,n}^+$ denote the set of all 1D triangulations such that all internal edges are not negatively oriented. From the well known map between lattice paths and 1D triangulations, see, e.g., [8], one has that $\Omega_{1,n}^+$ is in one-to-one correspondence with lattice paths (integer valued sequences with ± 1 increments) of length $2n$, which start at zero, end at zero, and stay non-negative. Moreover, under this correspondence, the total length

$L(\pi)$ of $\pi \in \Omega_{1,n}^+$ is precisely the area under the path π . The maximal element in $\Omega_{1,n}^+$, denoted $\hat{\pi}$, has total length

$$L(\hat{\pi}) = \sum_{i=1}^n i + \sum_{i=1}^{n-1} i = n^2.$$

Clearly, $\sigma \in \partial A$ implies that there exists a layer j such that $\sigma(j)$ has an internal vertical edge, i.e., an edge with horizontal length zero. In particular, on that layer we must have $L(\sigma(j)) \leq n^2 - n$. Therefore,

$$\begin{aligned} \frac{Z(\partial A)}{Z(A)} &\leq m \frac{\sum_{\pi \in \Omega_{1,n}^+: L(\pi) \leq n^2 - n} \lambda^{L(\pi)}}{\sum_{\pi' \in \Omega_{1,n}^+} \lambda^{L(\pi')}} \\ &\leq m \sum_{\pi \in \Omega_{1,n}^+: L(\pi) \leq n^2 - n} \lambda^{L(\pi) - n^2} \leq m \sum_{k \geq n} p_n(k) \lambda^{-k}, \end{aligned}$$

where we have estimated $\sum_{\pi' \in \Omega_{1,n}^+} \lambda^{L(\pi')} \geq \lambda^{L(\hat{\pi})} = \lambda^{n^2}$, and we use $p_n(k)$ to denote the number of $\pi \in \Omega_{1,n}^+$ which have $L(\pi) = n^2 - k$. From the correspondence with lattice paths it is seen that $p_n(k)$ is bounded by the number $p(k)$ of all partitions of the integer k . Using the well known fact that $p(k) = e^{O(\sqrt{k})}$ we arrive at the desired bound (6.18). \square

The lower bound in Theorem 6.1 can be improved as follows in the case $m \ll \sqrt{n}$.

Theorem 6.2. *For every $\lambda > 1$, there exists $c > 0$ such that the mixing time of the Glauber dynamics on triangulations of $\Lambda_{m,n}^0$, in the absence of constraints, satisfies $T_{\text{mix}} \geq \exp(cn^2/m)$ for all $n \geq m^2/c$, $m \geq 1$.*

Proof. Take $\varepsilon \in (0, 1)$ and divide the midpoints in $\Lambda := \Lambda_{m,n}$ into three regions as follows: Λ_ℓ is the set of $x \in \Lambda$ with horizontal coordinate between 0 and εn , Λ_c is the set of $x \in \Lambda$ with horizontal coordinate between εn and $(1 - \varepsilon)n$, and Λ_r is the remainder of Λ . Consider the set $\Lambda_1 \subseteq \Lambda_c$ of all internal midpoints $x \in \Lambda_c$ whose vertical coordinate is half-integer. Let $A \subseteq \Omega$ denote the set of triangulations σ such that every edge σ_x with $x \in \Lambda_1$ with vertical coordinate v is positively or negatively oriented according to the parity of $v + \frac{1}{2}$. In particular, adjacent layers are required to have one-dimensional edges of opposite orientation throughout the whole central region. This is a relaxed version of the set A appearing in the proof of Theorem 6.1 (see Figure 12), where the herringbone structure is now imposed only in the central region Λ_c .

We proceed by estimating $Z(\partial A)/Z(A)$ as in the proof of Theorem 6.1. Clearly, we can bound $Z(A)$ from below by considering only the maximal one-dimensional triangulations $\hat{\sigma}(j)$ in each layer j throughout the whole region Λ , and taking every edge with integer vertical coordinate to be horizontal (unit length), giving

$$Z(A) \geq \lambda^{mn^2 + O(mn)}.$$

Now let $\sigma \in \partial A$. The length of an edge whose midpoint is in $\Lambda_\ell \cup \Lambda_r$ is at most $\varepsilon n + m$. Since there are at most $8\varepsilon nm$ such midpoints, the total contribution of these edges to $|\sigma|$ is at most $8(\varepsilon n)^2 m + 8\varepsilon nm^2$. Moreover, there must be a midpoint $x_0 \in \Lambda_1$ whose edge σ_{x_0} is vertical (unit length). The contribution of the layer containing x_0 to the total length of σ is at most $(un)^2 + ((1 - u)n)^2$, where $\varepsilon \leq u \leq 1 - \varepsilon$. All other layers contribute total length at most $(m - 1)n^2$. Hence for any $\sigma \in \partial A$, we have

$$|\sigma| \leq 8(\varepsilon n)^2 m + 8\varepsilon nm^2 + (m - 1)n^2 + (u_* n)^2 + ((1 - u_*)n)^2,$$

where $u_* \in [\varepsilon, (1 - \varepsilon)]$ maximizes $(un)^2 + ((1 - u)n)^2$. Clearly, we have $u_* \sim \varepsilon$ (or by symmetry $u_* \sim (1 - \varepsilon)$). Since there are at most 2^{3mn} triangulations in total [1], we get

$$\frac{Z(\partial A)}{Z(A)} \leq \lambda^{8(\varepsilon n)^2 m + 8\varepsilon nm^2 - n^2 + (\varepsilon n)^2 + ((1-\varepsilon)n)^2 + O(mn)}.$$

Thus taking $\varepsilon = c/m$ and $n \geq m^2/c$ for some sufficiently small constant $c > 0$ (independent of n, m) gives

$$\frac{Z(\partial A)}{Z(A)} \leq e^{-c'n^2/m}.$$

This completes the proof. \square

Note: While the above bound is tight in the case $m = O(1)$, it becomes progressively weaker as m increases with n , and vacuous when $m \sim \sqrt{n}$. In particular, it is far from the conjectured behavior $T_{\text{mix}} = e^{\Omega(mn(m+n))}$ stated in Conjecture 1.1. The deterioration of the bound as m increases is apparently an artifact of our proof, which is based on an essentially one-dimensional ‘herringbone’ structure similar to that used in the proof of Theorem 6.1.

We end this section with a much weaker lower bound that holds for all $\lambda > 0$. This shows in particular that the upper bound of Theorem 5.1 for the small λ regime is tight.

Proposition 6.3. *There exists $c > 0$ such that the mixing time of the Glauber dynamics on triangulations of $\Lambda_{m,n}^0$ in the absence of constraints satisfies $T_{\text{mix}} \geq cmn(m+n)$ for all $m, n \geq 1$ and all $\lambda > 0$.*

Proof. Assume that $n \geq m$. Let $x = (1/2, n/2) \in \Lambda$ and let the initial condition σ be an arbitrary triangulation in which σ_x is the edge from $(0,0)$ to $(1,n)$. Let A be the set of triangulations in which σ_x is *not* positively oriented. By symmetry, the stationary probability of A is at least $1/2$ (with strict inequality if x is of type 1). Then, the total variation distance at time t can be bounded below as

$$\|p^t(\sigma, \cdot) - \mu\| \geq \frac{1}{2} - p^t(\sigma, A) \geq \frac{1}{2} - \mathbb{P}_\sigma(\tau_A \leq t),$$

where τ_A denotes the hitting time of the set A , and \mathbb{P}_σ is the probability on trajectories of the Markov chain with initial condition σ . Let $\Lambda' \subseteq \Lambda$ be the set of midpoints $(1/2, i)$ for $i = 0, 1/2, 1, \dots, n$. In order to flip σ_x all the way to an edge that is not positively oriented, we need to perform at least $\Omega(n^2)$ flips of the edges with midpoints in Λ' , as can be seen easily from the correspondence with lattice paths recalled in the proof of Theorem 6.1. Thus $\mathbb{P}_\sigma(\tau_A \leq t)$ is bounded above by the probability that out of t updates in Λ , at least $\Omega(n^2)$ of them are in Λ' . Each update independently falls in Λ' with probability $\frac{|\Lambda'|}{|\Lambda|} = O(1/m)$. By Markov’s inequality it follows that, for some constant $C > 0$,

$$\mathbb{P}_\sigma(\tau_A \leq t) \leq C \frac{t}{mn^2}.$$

Taking $t = mn^2/4C$ yields $T_{\text{mix}} = \Omega(mn(m+n))$. \square

7. FUTURE WORK

Our results suggest a number of immediate open questions.

- (1) Can the spatial mixing property (Theorem 1.2) and the polynomial mixing time bound (Theorem 1.3) be extended to the entire regime $\lambda < 1$? This would in particular complete the verification of part (a) of Conjecture 1.1. These questions, which are related, are likely to require a deeper investigation of the geometry of triangulations along the lines begun in Section 4.
- (2) Can the stronger exponential lower bound on mixing time in part (b) of Conjecture 1.1 be proved in the regime $\lambda > 1$? This will require the use of more sophisticated rigid triangulations than the “herringbone” structures used in Section 6.
- (3) The unweighted case $\lambda = 1$ seems particularly challenging, as our results strongly suggest that it corresponds to a “phase transition.” We conjecture that the spatial and temporal mixing properties here are likely to be very delicate.
- (4) In the “super-critical” regime $\lambda > 1$, the model appears to exhibit various “phases” according to the direction of alignment of the edges. It would be interesting to describe these phases and their contributions to the equilibrium distribution.

REFERENCES

- [1] E. E. Anclin, *An upper bound for the number of planar lattice triangulations*, Journal of Combinatorial Theory Series A **103** (2003), 383–386.
- [2] D. I. Dais, *Resolving 3-dimensional toric singularities*, Séminaires & Congrès **6** (2002), Soc. Math. France, 155–186.
- [3] J. DeLoera, J. Rambau, F. Santos, *Triangulations*, Springer Series on Algorithms and Computation in Mathematics Vol. 25, 2010.
- [4] J.-P. Eckmann, M. Younan, *Decay of correlations in a topological glass*, Philosophical Magazine, 2011, 1–22.
- [5] D. Eppstein, *Happy endings for flip graphs*, Proceedings of the 23rd Annual Symposium on Computational Geometry (SoCG), 2007, 92–101.
- [6] I. M. Gel’fand, M. M. Kapranov, A. V. Zelevinsky, *Discriminants, resultants and multidimensional determinants*, Birkhäuser, Boston, 1994.
- [7] S. Greenberg, A. Pascoe, D. Randall, *Sampling biased lattice configurations using exponential metrics*, Proc. 19th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 76–85 (2009).
- [8] V. Kaibel, G. Ziegler, *Counting lattice triangulations*, Surveys in Combinatorics, London Math. Soc. Lecture Note Ser. 307, pp. 277–307 (2003).
- [9] C. L. Lawson, *Transforming triangulations*, Discrete Mathematics **3** (1972), 365–372.
- [10] C. W. Lee, *Subdivisions and triangulations of polytopes*, in “Handbook on Discrete and Computational Geometry” (J. E. Goodman and J. O’Rourke, eds.), CRC Press, Boca Raton, FL, 1997, 271–290.
- [11] D. Levin, Y. Peres, E. Wilmer, *Markov Chains and Mixing Times*. AMS Bookstore, 2009.
- [12] F. Martinelli, *Lectures on Glauber dynamics for discrete spin models*, Lecture notes in Mathematics **1717**, Springer, Berlin, 1998, 93–191.
- [13] J. Matoušek, P. Valtr and E. Welzl, *On two encodings of lattice triangulations*, Manuscript, 2006.
- [14] L. McShine, P. Tetali, *On the mixing time of the triangulation walk and other Catalan structures*, DIMACS-AMS Volume on Randomization Methods in Algorithm Design **43** (1998), 147–160.
- [15] M. Molloy, B. Reed, W. Steiger, *On the mixing rate of the triangulation walk*, DIMACS-AMS Volume on Randomization Methods in Algorithm Design **43** (1998), 179–190.
- [16] S. Yu. Orevkov, *Asymptotic number of triangulations in \mathbb{Z}^2* , Journal of Combinatorial Theory Series A **86** (1999), 200–203.
- [17] J. Reinhard, K. Mecke, *Entropy of lattice triangulations*, Proceedings of the German Physics Society, Regensburg, 2010.
- [18] M. Sharir, E. Welzl, *Random triangulations of planar point sets*, Proceedings of ACM Symposium on Computational Geometry, 2006, 273–281.

- [19] B. Simon, *Statistical mechanics of lattice gases, Vol. I*, Princeton University Press, 1993.
- [20] B. Sturmfels, *Gröbner bases and convex polytopes*, University Lecture Series Vol. 8, American Mathematical Society, Providence, RI, 1995.
- [21] O. Ya. Viro, *Real plane algebraic curves: Constructions with controlled topology*, Leningrad Math. J. **1** (1990), 1059–1134.
- [22] E. Welzl, *The number of (and random) triangulations of planar point sets*, Presentation at Oberwolfach Workshop on Combinatorics, Probability and Computing, November 2006.
- [23] D. B. Wilson, *Mixing times of Lozenge tiling and card shuffling Markov chains*, Ann. Appl. Probab. **14** (2004), 274–325.
- [24] G. M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics **152**, Springer Verlag, New York, 1995.

PIETRO CAPUTO: DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ROMA TRE, LARGO SAN MURIALDO 1, 00146 ROMA, ITALY. EMAIL: caputo@mat.uniroma3.it

FABIO MARTINELLI: DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ROMA TRE, LARGO SAN MURIALDO 1, 00146 ROMA, ITALY. EMAIL: martin@mat.uniroma3.it

ALISTAIR SINCLAIR: COMPUTER SCIENCE DIVISION, UNIVERSITY OF CALIFORNIA, BERKELEY CA 94720-1776, U.S.A. EMAIL: sinclair@cs.berkeley.edu

ALEXANDRE STAUFFER: DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ROMA TRE, LARGO SAN MURIALDO 1, 00146 ROMA, ITALY. EMAIL: alexandrestauffer@gmail.com